Procurement Auctions with Variable Quantities

Dissertation

zur Erlangung des Grades eines Doktors der Wirtschaftswissenschaft

der Rechts- und Wirtschaftswissenschaftlichen Fakultät

der Universität Bayreuth

Vorgelegt

von

Claudio Rottner

aus

Karlsruhe

Dekan:	*Prof. Dr. Claas Christian Germelmann*
Erstberichterstatter:	*Prof. Dr. Stefan Seifert*
Zweitberichterstatter:	*Prof. Dr. Fabian Herweg*
Tag der mündlichen Prüfung:	*11.03.2026*

Acknowledgements

I am grateful to all those who helped me get here.
To those who paved the way,
and those who shaped my path in ways they may never know.
To those whose influence I recognise now,
and those whose contributions I will only see further down the road.

Abstract

In the procurement context, auction theory traditionally considers an auctioneer who buys a fixed number of goods from one (single-winner auctions) or multiple (multi-winner auctions) winning bidders. In these quantity auctions, the buyer purchases the fixed quantity from the lowest-cost bidder(s). This dissertation analyses procurement auctions, where the quantity depends on the bids of potential suppliers.

Chapter 2 considers a buyer who has quasi-linear preferences but suffers from buying too little (quantity goal) and spending too much (spending goal). They procure from one of several suppliers, each with private information on their constant unit costs. In the optimal mechanism, the buyer varies the procured quantity in the bids submitted by suppliers. However, if the goals are sufficiently demanding, the buyer also engages in bunching at their target quantity or their target spending; they procure their target quantity or spend exactly their budget for a range of bid prices. When one goal dominates all other concerns, the buyer engages in full bunching: A traditional single-winner quantity auction is optimal for a dominant quantity goal, and a single-winner budget auction is optimal when the spending goal is dominant.

Chapter 3 shows that a first-price single-winner auction with variable demand can implement the optimal mechanism even outside these extreme cases. However, these optimal auctions are demanding for the buyer in terms of their informational requirements and mental complexity. Addressing these concerns, Chapter 3 considers robust auctions that perform well regardless of the distribution of types, as well as simple (i.e., easy-to-implement) auctions. In a robust auction, the buyer maximises their expected utility for the worst possible distribution of cost types. Any demand schedule that secures the first-best quantity from the highest-cost type satisfies this criterion. Chapter 3 also compares the two simple setups of quantity and budget auctions. A buyer who does not suffer too much from variation in the procured quantity prefers the budget auction if both goals are equally important.

Chapter 4 compares quantity and budget auctions in a multi-winner setting. It investigates a multi-unit pay-as-bid budget auction where the auctioneer maximises the quantity procured with their predetermined, secret budget. Compared to quantity auctions, where the fixed traded quantity is unknown to the bidders,

budget auctions lower the auctioneer's costs by introducing an additional interaction between a bidders bids; bidders not only weigh a higher profit margin on a unit against a lower probability of supplying that unit; a higher margin on some unit also reduces the probability that the budget suffices to procure more units from the bidder.

Finally, Chapter 5 extends the analysis of multi-winner auctions with variable demand to electricity markets, where demand falls when prices increase without keeping total spending constant. Chapter 5 compares pay-as-bid and uniform pricing when symmetric producers with a weakly increasing marginal cost function compete for an uncertain, downward-sloping demand. Both price rules achieve first-best welfare when marginal costs are flat or there is an infinite number of producers. For a linear demand function with a uniformly distributed random intercept and linearly increasing marginal costs, pay-as-bid pricing results in a higher expected consumer surplus. In this linear model, pay-as-bid pricing yields higher expected welfare when demand variability is low but the overall welfare ranking remains ambiguous.

Contents

List of Figures

List of Tables

x

"[T]here is something inherently intriguing about an auction [...]. In the process of discussing this topic with hundreds of lay individuals over the years[,] [...] I have yet to find one whose eyes do not light up when the subject is discussed in depth."

— Cassady (1967/1980)

1 Introduction

Between his descriptions of the Greco-Persian Wars and an account of more-than-fox-sized, gold-digging ants,[1] the Greek historian Herodotus took the time to give a detailed description of a particular market institution in Ancient Babylon: the auctioning of future brides. According to Herodotus, the women to be married in a given year were sold to their future husbands by auction.[2] The future brides were auctioned consecutively in order of their beauty, beginning with the most beautiful one. The less desirable brides would be "assigned to whoever would take [them] with the least money"[3], i.e., bids would take the form of *asking* for less money for marrying the woman.

Herodotus' account is commonly considered to be the first description of an auction – and a complex one at that. Although such auctions almost certainly never took place (McNeal, 1988), Herodotus' description from the 5th century BC shows that people at the time were familiar with auctions. Whoever came up with the Babylonian bridal auctions also thought about strikingly modern aspects of

[1] Herodotus, *Histories* 3.102.
[2] Herodotus, *Histories*, 1.196.
[3] Ibid, in the translation of William Beloe, London: Leigh and Sotheby, 1791.

what we today call *market design*. For example, the Babylonian bridal auctions moved directly from a forward to a reverse auction: An auction starts by selling brides and ends with using (parts of) the revenue to buy husbands. A similar idea has recently gained some popularity in spectrum auctions. For example, the Austrian government used part of the revenue it made from selling rights to use frequencies for mobile networks to telecommunication companies to buy coverage for remote areas from the very same companies in an "innovative auction design"[4] in 2020. Thinking about such details of designing an auction, of course, takes for granted that buying and selling by auction is a good idea to begin with. But why would someone sell by auction instead of posting a price like in a supermarket or on Amazon?

1.1 A Brief History of Auctions and Mechanism Design

Even though there is a long history of selling more or less important goods by auction, from the whole Roman empire in 193 AD (Cassady, 1967/1980, p. 29) to Justin Timberlake's half-eaten French toast in 2006 (Mackay, 2017), we do not know much about why people started using auctions. Anecdotally, it sounds plausible that auctions could have developed when sellers faced high demand on markets: Eager to obtain a scarce good, buyers might have started shouting out what they were willing to pay to prevent the seller from giving the good to someone else at a lower price. Today's famous flower auctions in the Netherlands have emerged exactly this way. They have their origin in 1887, when a vegetable grower named Jongerling successfully sold his produce by auction in a high-demand situation (Cassady, 1967/1980, p. 37). This account makes it seem as if the primary role of auctions was to spur competition. However, their real purpose is to overcome *asymmetric information*. Had Jongerling known what each potential buyer was

[4] That is, according to the regulatory authorities overseeing the auction (see RTR, 2019).

willing to pay, there would have been no need to have buyers compete in an auction. He could have offered the good to the buyer with the highest willingness to pay at that exact price and would have realised the highest possible profit. Auctions leverage competition between buyers, i.e., the risk of not obtaining the good, to prevent buyers from understating their true willingness to pay too much.

In the 1970s, economists discovered that the analysis of asymmetric information is a crucial field for economic analysis. Asymmetric information is the explanation for the existence of many (market) practices, from insurance premiums (Rothschild and Stiglitz, 1976), over (potentially unproductive) education (Spence, 1973) to – as we have seen – auctions. By the same token, markets where profitable trade is possible can break down if such practices do not develop, as the seminal paper by Akerlof (1970) revealed: Nobody in the market for a high-quality used car would ever buy a used car unless used car-dealers (the not sketchy kind) find a way to signal to their uniformed buyers that the cars they are selling are genuinely in good condition and not only look as if they were.

If markets with asymmetric information developed "additional" practices beyond posting prices, such as varying insurance premiums or signaling of quality, it is natural to ask what the optimal practices (for whoever gets to decide on them) are. This is the question *mechanism design* tries to answer, where a *mechanism* is usually simply a rule for allocating goods. Given the sheer number of ways one could potentially allocate goods – for example, goods could go to the party writing the best essay on their moral desert of the good – it required two important insights to make headway. First, Hurwicz (1972) introduced the idea of *incentive compatibility*. Designing a good mechanism requires that market participants reveal (parts of) their *private information*, i.e., information that only they have. Because the designer of the mechanism can (generally) not force them to do so, a meaningful mechanism must give them incentives to report their private information to the designer. For example, consumers voluntarily reveal private information on the riskiness of their behaviour when choosing an insurance with a particular premium, bidders reveal information on their value for a good when placing a bid in an auction, potential employees signal their intelligence through education, and used car dealers reveal information on the quality of cars through third-party certifications. Second, a string of papers (among others, Gibbard, 1973; Myerson, 1979, 1982, 1986) developed ever more general versions of what is nowadays known as the *revelation principle*: We can safely ignore the essay-writing mechanisms. Any outcome that they might produce can also be reached by a mechanism in which market participants directly pass on the information they have, with the mechanism designer incentivising them to do so truthfully. With the revelation principle in place, the problem of finding optimal mechanisms became a mathematically much more tractable problem, which finally allowed economists to give a formal rationale for using auctions: Myerson (1981) showed that auctions – although they generally do not involve bidders truthfully revealing their maximum willingness to pay – maximise the auctioneer's expected profits (under some regularity conditions). Thus, two and a half millennia after Herodotus reported on auctions, economists finally had a convincing reason for their use.

However, the result of Myerson (1981) only established the optimality of *single-unit auctions*. Single-unit auctions trade only one good, or rather, a single bundle of goods. Consequently, only one bidder can win the auction, which is why these auctions would be more aptly called single-winner auctions. Most auctions with high real-world relevance – from spectrum auctions to the sale of treasury bills and trade on electricity markets – are multi-winner auctions; there is more than one winning bidder and the auction splits the goods among them. The same even holds for the Babylonian bride auctions. Consequently, much of what is referred to as *auction theory* focuses on the analysis of these multi-winner auctions. In contrast to single-winner auctions, most widely-used multi-winner auctions are not generally optimal. The closest to an optimal multi-winner auction is the Vickrey-Clarke-Groves auction[5] (Clarke, 1971; Groves, 1973), which maximises welfare, i.e., the sum of profits over both bidders and the auctioneer. However, as the auctioneer often aims to maximise their own profits rather than welfare (which attaches value to the profits of bidders), this property of the Vickrey-Clarke-Groves auction does not qualify it as an optimal auction in most settings either.[6] Thus, as was the case throughout most of history, real-world economic agents seem to believe that auctions are a smart way of trading goods in light of private information even though there is no reason to believe that they are the optimal way to do so. The upshot is that much of the auction literature analyses multi-winner auctions that are of high practical relevance, but generally not optimal. This ambiguous relation of auction theory to optimal mechanism design is as old as auction theory itself: The foundational works of auction theory (Vickrey, 1961, 1962) describe optimal bidding behaviour in single-winner auctions purely for its practical relevance. The connection to optimal mechanisms only came twenty years later.

1.2 Procurement Auctions with Variable Quantities

For most people, the mental picture they associate with an auction is the sale of valuable works of art or antiquities. Besides also considering multi-winner auctions, this dissertation departs from this mental picture in two ways. First, it considers procurement instead of sales auctions: the auctioneer buys from the bidder(s) offering the lowest price rather than selling to those bidding the highest price(s). Conceptually, the difference is negligible. Procurement auctions are

[5]This auction allocates goods in a way that maximises the bidders' aggregated (reported) willingness to pay for the bundles of goods they receive. The price a bidder pays is the (reported) loss other bidders have from not obtaining the goods that go to the bidder.

[6]Linnenbrink (2025) shows that the pay-as-bid auction can generate higher profits for the auctioneer than the Vickrey-Clarke-Groves auction.

no different from sale auctions with reversed signs. Second, this dissertation analyses auctions with variable rather than fixed quantities. When auctioning a unique piece of art, the number of items up for sale is fixed by construction – the whole value of the item is that there is only one. Similarly, there was a fixed number of brides in the Babylonian bride auctions. Whatever the way the (hypothetical) Babylonians determined the number of brides in a given year, this number was, apparently, independent of what happened at the auction. However, this example also illustrates that there is something strange about that. The Babylonian auctioneer paid to get all brides married rather than stopping the auction when nobody was willing to pay a positive price. From an economic point of view, it would seem more reasonable to reduce the number of auctioned brides in response to the observed bids and only sell those brides that fetch a positive price and, thus, generate profit to the auctioneer. In standard economic terms, this corresponds to the idea that there is a *supply curve*, which describes how (profitable) supply changes with the market price. In procurement auctions, the equivalent is a downward-sloping demand by the auctioneer; they curb their demand when the prices in the auction are high. Usually, such behaviour is considered "normal" in economics. Economists sometimes even call the inverse relationship between price and demand the "law of demand" as if it were a law of nature. Thus, the idea of fixing the traded quantity is somewhat a peculiarity of auction theory.

Are there procurement auctions with variable demand in the real world? This dissertation analyses two examples of such auctions: electricity markets (Chapter 5) and the fourth stage of the Austrian spectrum auction in 2020, in which the Austrian government bought coverage for undersupplied communities (Chapter 4). In electricity markets, energy suppliers bid on satisfying a demand that goes down when prices increase. In the Austrian coverage auction, the government reduced demand in response to price increases in such a way that they would spend the same amount no matter the prices. In the terminology of this dissertation, they conducted a *budget auction*; they fixed a budget for procurement and sourced as many units as possible with their budget.

Although there is a limited literature on this kind of auction in the single-winner context (Dastidar, 2008; Deck and Wilson, 2008; Liu, 2007; Liu and Parlour, 2023), the topic still seems underexplored given that organisations of all forms – from governments to companies – regularly fix budgets for specific projects to keep spending under control. This practice is so pervasive that even consumers have budgets or "mental accounts" for their spending in different categories, such as clothes, food or vacations and do not readily move money between these mental accounts (Thaler, 1985, 1999).[7] It seems only natural to extend this logic to procurement auctions such that the buyer knows they will keep within their budget.

1.3 The Contribution of this Dissertation

This dissertation consists of four single-authored chapters on procurement auctions with a variable traded quantity. It places a particular focus on budget auctions. First, Chapter 2 derives an optimal procurement mechanism when the buyer responds to prices but also has a quantity and a spending goal. It shows that procurement auctions with a fixed quantity and a fixed budget implement this optimal mechanism when the respective goal is sufficiently important to the buyer. Chapter 3 demonstrates that auctions with a variable demand schedule can implement the optimal mechanism more generally. Moreover, the chapter finds that a buyer who attaches equal weight to their quantity and spending goal, by tendency, prefers a budget over a quantity auction because the average unit they buy is cheaper in the budget auction. Chapter 4 provides an analogous result for

[7]Maybe the direction of causality is also the other way round, and organisations behave this way because thinking in budgets is inherently human.

the multi-winner context: When each unit trades at the price bid on that specific unit, bids are lower in a budget auction than in a comparable quantity auction. Chapter 5 considers a similar model with a variable demand schedule that does not fix spending to address one of the key questions in practical market design for electricity markets: How should prices be determined?

Chapter 2: Unifying Perspectives on Optimal Procurement in the Single-Winner Context.— For the single-winner context, the previous literature has established that budget auctions are optimal under the same conditions as quantity auctions: Traditional quantity auctions minimise the expected spending required to buy a fixed quantity when the suppliers have private information on their constant unit costs and it is common knowledge that these unit costs are independent draws from some known distribution. Under the same conditions, budget auctions maximise the quantity a buyer can procure with a fixed budget (Liu, 2007). Moreover, the literature on mechanism design has established the optimal procurement mechanism when the buyer has quasi-linear preferences (Dasgupta and Spulber, 1989).

Chapter 2 unifies these perspectives. They are special cases of the optimal mechanism in a model where the buyer has quasi-linear preferences but suffers when buying too little or spending too much. When these goals – the minimum quantity the buyer wants to source and the maximum spending they are willing to accept – are irrelevant, the buyer trades off the quantity they buy with the marginal price they pay. Every increase in the price results in a decrease in demand and, generally, a change in total spending. The quantity and spending goals represent "sticking points" in this trade-off: As the buyer wants to buy a minimal quantity and spend less than their budget, they will keep the procured quantity or total expenditures at their target level, even if there is an infinitesimally small change in the marginal price. Consequently, a sufficiently demanding and important quantity goal implies that the buyer procures a fixed quantity – the traditional setting of auction theory. Analogously, if the buyer's budget is very tight and overspending is very painful to them, they fix their total expenditures independently of the price they pay — the setting of budget auctions. When both kinds of goals are relevant, the optimal mechanism will generally see a price range for which the buyer sources the same quantity and a price range for which the buyer keeps their total expenditures constant.

Chapter 2 adopts a traditional mechanism design perspective, focusing on optimal direct revelation mechanisms, i.e., suppliers directly, truthfully reporting their costs. For the special cases of fixed-quantity and fixed-budget procurement, it is clear from previous literature that auctions implement this optimal mechanism.

Chapter 3: Single-Winner Auctions as Procurement Mechanisms.— Chapter 3 explores the connection between the optimal mechanism and auctions outside these special cases. It shows that first-price auctions, where the winning bidder gets paid their ask price, can implement the optimal mechanism more generally. All this requires is that the traded quantity is determined from the winning bid via a pre-announced demand schedule, which the buyer can compute from the known distribution of suppliers' costs. Although this is a satisfactory result from a theoretical perspective, it does not seem like many real-world auctioneers conduct such auctions; the feeling creeps back that auctions are about something more practical than optimality. Chapter 3 considers two explanations for why real-world buyers might not use optimal auctions. First, they might opt for "robust" auctions, which do not require any knowledge of the distribution of suppliers' costs. According to the commonly used maxmin expected utility criterion, the buyer should then choose the auction that does best for the worst-case distribution. Any auction that results in the buyer sourcing the first-best quantity – i.e., the quantity they would buy if they knew the winning supplier's costs – when the winning bidder has the highest imaginable costs, is maxmin-optimal. Second, real-world auctioneers might use "simple" auctions to avoid mental complexity, such as computing the optimal demand schedule. Chapter 3 compares two auction designs that are intuitively simple: fixed-quantity and budget auctions, both of which are optimal if the respective goal is sufficiently important. Chapter 3 reveals that a buyer who does not suffer too much from *variation* in the procured quantity prefers the budget auction even if they attach slightly more weight to their quantity goal. The reason is that the budget auction allows them to buy a higher expected quantity with the same expected spending: In a budget auction, the buyer buys more when unit costs are low and less when unit prices are high. Further strengthening this result, suppliers submit lower bids in the budget auction: By lowering their bids, they not only increase their probability of winning the auction (as in a quantity auction) but they also increase the number of (profitable) units that they supply.

Chapter 4: Fixed-Quantity and Budget Auctions in the Multi-Winner Setting.— Something similar also holds in a multi-winner setting. Chapter 4 considers a multi-winner budget auction where the buyer auctions a fixed but secret budget. The bidders submit bid functions that specify how much money they demand for supplying a given quantity. The buyer chooses the combination of different bidders' bids that maximises the quantity they source with their budget, when they pay the respective bidder's ask price for each unit (pay-as-bid). The bidders do not know the buyer's budget but have some probabilistic belief about it. Compared to a quantity auction, where the buyer would fix a secret traded quantity and minimise the total spending required to reach that quantity, bids in the budget auction are lower. In the budget auction, bidders run the risk that the budget does not suffice to source more (profitable) units from them when they increase their bids. As in the single-winner case, they lower their bids to increase the – in this case, expected – quantity of profitable units they supply.

Chapter 4 assumes that costs are common knowledge among bidders. This is a common assumption in much of the relevant literature (especially, Klemperer and Meyer, 1989; Pycia and Woodward, 2026), as deriving optimal bids with unknown costs of other bidders has proven to be difficult for multi-winner auctions. However, without any uncertainty for bidders, there is no meaningful equilibrium in a pay-as-bid auction, i.e., no matter how the bidders bid, there is always at least one bidder who should change their bidding strategy (see Chapter 4.5). Consequently, the literature, including this dissertation, replaces uncertainty about competitors' costs with uncertainty about demand.

While this set-up may seem artificial, there are applications that mirror the combination of known costs and uncertain demand very well. For example, the model in Chapter 4 is not only consistent with previous literature, but is also a fairly accurate description of the fourth stage of the Austrian spectrum auction in 2020: To prevent collusion, the government kept its budget for procuring network coverage for undersupplied communities secret, introducing demand uncertainty for bidders. At the same time, the auction design ensured that bidders had similar cost structures: each bidder selected the communities for which they would provide coverage from a list that was specific to the bidder. The regulatory authority designed these lists to equalise costs across bidders. Even though bidders' costs may not have been perfectly identical, it is unlikely that bidders incurred the costs to find out to what degree the others' costs might diverge, making the assumption of commonly known costs a reasonable approximation.

Chapter 5: Pay-As-Bid vs. Uniform-Price Auctions in Electricity Markets.— Chapter 5 considers another real-world setting, where the assumptions of uncertain demand and commonly known costs are likely to hold: electricity markets. As all electricity suppliers use the same technologies for electricity generation and the power plants a company owns are generally observable to competitors, electricity suppliers should have reliable estimates of their competitors' costs. Moreover, they face uncertainty over demand as bidding takes place before demand realises: In electricity markets, demand and supply must be equal at every moment. Consequently, there is no time to observe demand before suppliers bid on serving that demand. As a result, electricity suppliers submit their bids before demand is fully known. Analysing this market, therefore, only requires minor tweaks to the model from Chapter 4.

Chapter 5 adapts the model from Chapter 4 to compare two pricing rules in electricity markets. In electricity markets, there is a stop-out price that equalises supply and demand. Suppliers are willing to offer as many "units" of electricity for weakly less than the stop-out price as there are "units" of electricity for which buyers are willing to pay weakly more than the stop-out price. Most electricity markets today are organised as uniform-price auctions: All units trade at the stop-out price. By contrast, in a pay-as-bid auction, suppliers only get their ask price, which is weakly lower than the stop-out price. Chapter 5 finds that both auction designs achieve the highest possible social welfare when marginal costs are flat or there is an infinite number of producers. For a linear demand function with a uniformly distributed random intercept and linearly increasing marginal costs, Chapter 5 shows that consumers are better off in a pay-as-bid auction. Also accounting for the value of suppliers' profits, the pay-as-bid auction generates more value for society if demand does not vary too much.

2

Optimal Procurement between Quantity and Spending Goals

Unpublished Working Paper

Abstract

The literature on optimal procurement usually considers a buyer with quasi-linear utility; the buyer determines the quantity they source based on the marginal price. By contrast, auction theory shows that auctions are optimal when the buyer has a clear goal that is independent of the marginal price: An auction minimises the spending required to source a fixed quantity or maximises the procurement with a fixed budget. This chapter considers optimal procurement with a generalised preference structure. The buyer has quasi-linear utility but suffers when procuring too little or spending too much. They procure from one of several suppliers, each with private information on their constant unit costs that are independent draws from a commonly known distribution. This chapter has three core findings. First, introducing relevant goals leads to bunching in the respective dimension: The buyer procures their target quantity or spends exactly their budget for a range of supplier types. Second, as the goals become more important, the size of these bunching intervals increases. When one goal dominates all other concerns, the buyer engages in full bunching; auctions fixing the quantity or total spending are optimal. Third, quantity and spending goals differ; spending goals affect the buyer's demand even when the buyer meets their goal. With a spending goal, the buyer curbs their demand even when they do not overspend, as this reduces the information rents of more efficient suppliers; this lowers spending on more efficient types and, therefore, reduces the penalty for overspending when contracting with these types.

2.1 Introduction

In 2024, the British Government announced it would auction off a budget worth 1.5 billion British pounds to support renewable energy (UK Government, 2024). Similarly, the European Commission provided a fixed budget of 800 million euros for its auction to support hydrogen production in 2024 (European Commission, 2024). The Austrian telecommunications regulatory authorities used a "fixed-budget auction" to procure network coverage for remote areas in 2020 (RTR, 2020).

The preferences underlying this type of fixed-budget procurement mirror those traditionally posited in auction theory. The buyer fixes total spending, or more traditionally, the procured quantity, and adheres to this goal no matter the marginal price of one unit. In the fixed-quantity case, buyers do not lower the quantity they buy, even if this could substantially lower total costs. In the fixed-budget case, they spend exactly their budget, even if an additional unit is almost costless.

As these assumptions are quite extreme, much of the literature on optimal procurement models a buyer with quasi-linear preferences (e.g., Dasgupta and Spulber, 1989; Laffont and Martimort, 2002); the buyer has a positive but decreasing consumption utility and total expenditures lower the buyer's utility as the numeraire. Thus, the buyer only procures an additional unit if its marginal utility exceeds its marginal price.

In all of these cases, the structure of optimal procurement is well-known. Whether spending a fixed budget or procuring a fixed quantity, a simple auction implements the optimal mechanism when sourcing from one supplier. To source a fixed quantity at a minimal expected cost, the buyer should procure the entire quantity from the supplier who commits to providing that quantity at the lowest cost in a fixed-quantity auction (Myerson, 1981). To buy the maximum amount with a given budget, the buyer should conduct a budget auction (Liu, 2007), i.e., award the entire budget to the supplier who commits to providing the highest quantity in exchange. With quasi-linear preferences, the buyer procures more when the unit costs of the winning supplier are lower (e.g., Dasgupta and Spulber, 1989), which, generically, translates to payments differing in the type of the winning supplier.

Regarding preferences, much of actual procurement seems to fall between these cases: In contrast to the quasi-linear model, specific targets do matter. For example, companies must procure the right quantity of crucial inputs. Similarly, government projects need to run within their predefined budgets. However, the marginal price of procurement also matters, which is ignored by the "auction-style" preferences. What is more, quantity and spending goals can coexist. Consider a government that wants to procure additional capacity for (subsidised, privately-owned) renewable electricity production to safeguard the national supply and meet climate targets – a specific quantity goal. Simultaneously, the government might want to limit the corresponding spending to avoid debt financing for ideological (e.g., government budgets should be balanced) or economic (e.g., debt financing requires interest payments) reasons. Either way, exceeding a spending target results in some sort of disutility.

This chapter aims to bridge the gap between the traditional auction literature, which posits unalterable quantity or spending goals, and the literature on optimal procurement that considers the trade-off between procured quantity and spending but ignores the importance of specific goals. The contribution of this chapter is to derive the optimal procurement mechanism for a generalised preference structure. The buyer has quasi-linear preferences but also (to a varying degree) suffers from procuring too little and spending too much, which, formally, mirrors models of loss aversion (e.g., Kőszegi and Rabin, 2006, 2007, 2009; Lange and Ratan, 2010; Shalev, 2000). Quasi-linear preferences, the preference structure commonly imposed in auction theory, and a loss-averse buyer are nested in this preference structure. The potential suppliers have constant unit costs representing independent draws from a commonly known distribution; the buyer sources from (at most) one supplier (SIPV model).

This chapter documents three features of optimal procurement with quantity and spending goals. First, introducing a relevant quantity or spending goal results in partial bunching in the quantity (Section 2.3.2) or spending (Section 2.3.3), respectively: There is a subset of adjacent (cost) types from which the buyer sources exactly their target quantity or on which they spend their target budget. If the buyer has both a quantity and a spending goal (and neither dominates all other concerns), the optimal mechanism generally results in partial bunching regarding quantity *and* spending (Proposition 2.1 / Section 2.3.4). Second, as meeting a goal becomes more important, the length of the corresponding bunching interval increases. Thus, if the quantity or the spending goal dominates all other concerns, the buyer fixes the procured quantity or total expenditures, respectively. Third, there is a difference between quantity and spending goals. With a quantity goal, the buyer boosts their demand when sourcing from types for which they would fall short of the desired quantity when ignoring the goal. With a spending goal, they curb their demand even when contracting with types for which they would keep within their budget anyways. By doing so, they reduce the information rents of more efficient types, which lowers the (potential) penalty for overspending on the more efficient types.

These findings link existing literature on optimal procurement with quasi-linear preferences and the literature on auctions. When the goals are irrelevant, the optimal mechanism corresponds to that for quasi-linear utility (e.g., Dasgupta and Spulber, 1989) without any bunching. The first finding of this chapter is that introducing relevant goals results in bunching because of the kink in the buyer's utility function. Consider the case of a quantity goal. If the procured quantity

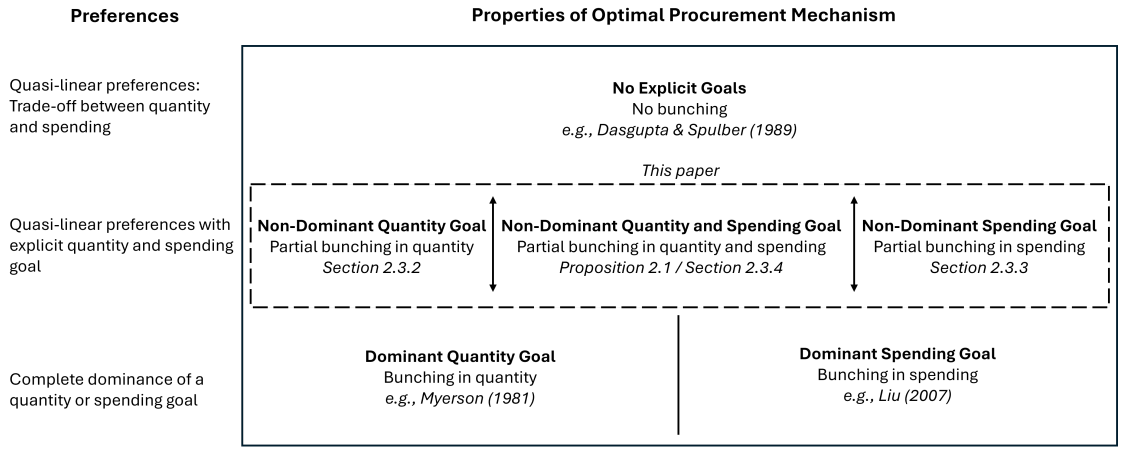

Figure 2.1: Optimal mechanism for a buyer who has quasi-linear preferences and suffers from missing their quantity or spending target and its relation to previous literature.

falls short of the goal, the marginal utility of procuring more jumps up; buying a higher quantity now helps avoid the penalty for missing the goal. Thus, there can be a range of supplier types for which the buyer's marginal utility of sourcing an additional unit is positive if they procure less than "prescribed" by their goal but is negative if they procure more than their target quantity. This results in bunching at the target quantity. The same logic applies to a spending goal. As the goals become more important, the jump in the marginal utility and, hence, the size of the bunching intervals increases. Consequently, when one goal dominates the buyer's utility, the buyer engages in full bunching. This full bunching can be implemented by a (first-price) quantity, or budget auction, which conforms with well-known results from auction theory. The optimality of standard auctions to procure a fixed quantity or procurement with a fixed budget has been established by Myerson (1981) and Liu (2007), respectively. Bidding behaviour in fixed-quantity auctions is well understood (e.g., Krishna, 2010). Dastidar (2008), Deck and Wilson (2008), and Liu and Parlour (2023) consider bidding behaviour in fixed-budget auctions. Figure 2.1 summarises the core findings on the optimal mechanism and their relation to previous literature.

This chapter proceeds as follows. Section 2.2 presents the model setup. Section 2.3 derives the optimal procurement mechanism. Section 2.4 discusses the findings before Section 2.5 concludes.

2.2 Model Setup

2.2.1 Buyer's Utility between Quantity and Spending Goals

The buyer evaluates the purchase of quantity x at total cost t when the corresponding goals are $\hat{x}$ and $\hat{t}$, according to the utility function

$$u(x,t) = \underbrace{\Big(v(x) - t\Big)}_{\text{Direct utility}} \underbrace{-\lambda_Q \, \mathbf{1}_{x \leq \hat{x}} \, (\hat{x} - x) - \lambda_B \, \mathbf{1}_{t \geq \hat{t}} \, (t - \hat{t})}_{\text{Penalty functions: penalty for missing goals}}, \tag{2.1}$$

where $\lambda_Q, \lambda_B \geq 0$ reflect the importance of hitting the quantity (subscript Q as in "quantity") and spending goal (subscript B as in "budget"), respectively; $v(x)$ is an increasing and concave function with $\lim_{x \to 0^+} v'(x) = \infty$ and $\lim_{x \to \infty} v'(x) = 0$. The indicator functions take the value one if the procured quantity (total spending) is weakly below (weakly exceeds) its target $\hat{x}$ ($\hat{t}$) and are zero otherwise. For compactness, denote them by $\mathbf{1}_Q$ ($\mathbf{1}_B$) in the following. Thus, the buyer suffers when buying less than $\hat{x}$ or spending more than $\hat{t}$. For simplicity, (2.1) assumes that these penalties are linear in the deviation from the target. The core results of this chapter hold as long as the penalty functions are weakly convex (see Section 2.4), which includes measuring the penalty for missing the quantity goal in terms of utility, i.e., by $v(\hat{x}) - v(x)$. Due to the quasi-linear direct utility, the buyer generally makes a trade-off between the procured quantity and total spending. The penalty functions give importance to reaching specific goals.

Fixed-Quantity and Fixed-Budget Procurement as Limiting Cases.— The utility function in (2.1) nests the case of quasi-linear preferences ($\lambda_Q = \lambda_B = 0$) as well as the fixed-budget and the fixed-quantity case. For $\lambda_B \to \infty$, the penalty function enforces a strict upper bound on spending. However, the buyer only exhausts their budget if the consumption bundle that maximises the direct quasi-linear utility results in total expenditure above $\hat{t}$. Similarly, for $\lambda_Q \to \infty$, the penalty function enforces a strict lower bound on the procured quantity. The buyer will procure more than the target quantity if this maximises their direct utility. Thus, the fixed-budget and fixed-quantity cases correspond to situations where deviations from the respective goals are very painful to the buyer (high $\lambda_{B/Q}$), and the goals are demanding enough for the penalty functions to be relevant. Intuitively, this correctly captures the rationale for fixed-budget and fixed-quantity procurement: *Fixing* spending is not only the consequence of having a budget cap but a budget cap that is too tight to ignore. Similarly, fixing the procured quantity is best understood as a binding minimum quantity requirement.

Hence, the utility function in (2.1) incorporates the case of quasi-linear preferences and gives an intuitive interpretation of fixed-quantity and fixed-budget procurement. On the downside, the correspondence to the literature on fixed-quantity and fixed-budget auctions is imperfect. Unlike in this literature (e.g., Myerson, 1981 and Liu, 2007, respectively), the buyer is not a maximiser of the expected procured quantity or a minimiser of expected spending in the limiting cases. This has two important implications: First, due to the infinite slope of $v(x)$ at $x = 0$, the buyer will not accept zero-trade. Section 2.3 shows that this results in their optimal procurement strategy as a fixed-quantity/fixed-budget auction *without* a reserve price. By contrast, the standard results on the optimality of auctions require the buyer to set a reserve price. Second, as $u(\cdot)$ is (weakly) concave in x and t, the buyer (weakly) prefers a first- over a second-price auction in either of the extreme cases (Remark 2.1), whereas "normally" they are indifferent.

Behavioural Interpretation of the Model.— It is also possible to interpret the preferences from (2.1) as two-dimensional loss-aversion, where the buyer suffers from falling short of some psychological reference point in terms of the procured quantity and total spending. Formally, the buyer's utility in (2.1) is consistent with the behavioural literature: Just as in the seminal papers by Kőszegi and Rabin (2006, 2007, 2009) utility is the sum of a standard direct consumption utility and a reference-dependent component; in (2.1), the reference-dependent component is linear and only considers reference-dependent losses (see, for example, Lange and Ratan, 2010; Shalev, 2000, for similar specifications). Therefore, this chapter also adds to the literature using concepts from behavioural economics to understand organisational decision-making (e.g., Baucells *et al.*, 2024; Camerer and Malmendier, 2007; Gao and Guo, 2024; Long *et al.*, 2020). Regarding auctions, the literature has primarily focused on behavioural bidders (e.g., Balzer and Rosato, 2025; Balzer *et al.*, 2022; Filiz-Ozbay and Ozbay, 2007; Lange and Ratan, 2010; Rosato and Tymula, 2019; Shunda, 2009), whereas this chapter models a behavioural principal.

2.2.2 Potential Suppliers and the Structure of the Mechanism

Suppliers.— Regarding suppliers, this chapter considers a canonical SIPV model (e.g., Krishna, 2010): There are n potential suppliers. Each supplier i has constant unit costs θ_i, that are private information of supplier i. It is common knowledge that the suppliers' unit costs are independent draws from the distribution $F(\theta)$ on $[\underline{\theta}, \bar{\theta}]$ with $0 < \underline{\theta} < \bar{\theta}$ and $\frac{d}{d\theta}\frac{F(\theta)}{f(\theta)} \geq 0$ (monotone reverse hazard rate property), where $f(\theta) = F'(\theta)$. Suppliers maximise their expected profits. A supplier of type θ with a revenue of t has profits of $t - \theta x$ when they produce x units. Their outside option has a value of zero.

Structure of the Procurement Mechanism.— Mirroring some usual features of actual procurement (and similar to Li and Lo, 2023), this chapter only considers mechanisms that satisfy the following properties:

- **Lowest Reported Costs Win (LCW):** The buyer only procures from the supplier who claims to have the lowest unit costs.

- **Payment to Winner Only (PtW):** The buyer only pays the supplier from whom they source.

As the buyer evaluates the quantity they actually source and their actual expenditures relative to their goals, it is useful to represent a mechanism by a quantity schedule $x(\cdot)$ and a transfer schedule $t(\cdot)$ that specify the quantity sourced and the transfer paid to a supplier of a given type *conditional on the supplier winning the contract.* Under the assumptions (LCW) and (PtW), Remark 2.1 establishes that $x(\cdot)$ and $t(\cdot)$ only depend on the type of the winning supplier in the optimal mechanism, i.e., an optimal mechanism can be represented by $(x(\theta), t(\theta))$, where θ is the winning supplier's revealed type.

Remark 2.1 (Optimal Mechanism only depends on Type of the Winning Bidder) *Consider a mechanism* $(\tilde{x}(\theta, \boldsymbol{\theta}_{-i}), \tilde{t}(\theta, \boldsymbol{\theta}_{-i}))$ *where* θ *is the type of supplier* i *who wins the contract and* $\boldsymbol{\theta}_{-i}$ *is the vector of types of suppliers* $j \neq i$*. Assume that the mechanism* $(\tilde{x}(\cdot), \tilde{t}(\cdot))$ *induces the risk-neutral suppliers to reveal their true type and guarantees weakly positive expected profits to all supplier types. Given (LCW) and (PtW) and assuming that* $\tilde{x}(\cdot)$ *and* $\tilde{t}(\cdot)$ *vary in* $\boldsymbol{\theta}_{-i}$*, a buyer with the preferences from* (2.1) *strictly prefers the mechanism* $(x(\theta), t(\theta))$*, where* $x(\theta) = \mathbb{E}_{\boldsymbol{\theta}_{-i}}[\tilde{x}(\theta, \boldsymbol{\theta}_{-i})|\theta < \theta_j]$ *and* $t(\theta) = \mathbb{E}_{\boldsymbol{\theta}_{-i}}[\tilde{t}(\theta, \boldsymbol{\theta}_{-i})|\theta < \theta_j]$*. Every supplier reports the same type and has the same expected profits under either mechanism.*

Proof. Given that supplier types are independent, the expected profits of a supplier of type θ_i are identical under both mechanisms for every type $\tilde{\theta}$ they could report. The buyer's expected utility under $(\tilde{x}(\cdot), \tilde{t}(\cdot))$ is

$$\mathbb{E}[u(\tilde{x}(\cdot), \tilde{t}(\cdot))] = \int_{\underline{\theta}}^{\bar{\theta}} \mathbb{E}_{\boldsymbol{\theta}_{-i}} \left[u\Big(\tilde{x}(\theta, \boldsymbol{\theta}_{-i}), \tilde{t}(\theta, \boldsymbol{\theta}_{-i})\Big) \Big| \theta < \theta_j \right] \, f_1(\theta) \, d\theta,$$

where $f_1(\cdot)$ denotes the density of the distribution of the most efficient supplier's type. It follows from Jensen's inequality that

$$\begin{aligned}\mathbb{E}[u(\tilde{x}(\cdot),\tilde{t}(\cdot))] &< \int_{\underline{\theta}}^{\bar{\theta}} u\Big(\underbrace{\mathbb{E}_{\boldsymbol{\theta}_{-i}}[\tilde{x}(\theta,\boldsymbol{\theta}_{-i})|\theta<\theta_j]}_{=x(\theta)},\underbrace{\mathbb{E}_{\boldsymbol{\theta}_{-i}}[\tilde{t}(\theta,\boldsymbol{\theta}_{-i})|\theta<\theta_j]}_{=t(\theta)}\Big)\, f_1(\theta)\; d\theta \\ &= \mathbb{E}[u(x(\theta),t(\theta))]\end{aligned}$$

because the buyer's utility is strictly concave in x and weakly concave in t. Thus, the inequality is strict regarding the change of the quantity schedule and weak regarding the transfer schedule. ■

The concavity of $u(\cdot)$ implies that the buyer dislikes any variance of outcomes. Thus, they eliminate all variance in outcomes that does not create incentives for suppliers. Per assumption, the buyer only influences the pay-offs of the winning supplier (LCW and PtW). Because the winning supplier only knows their own type and, therefore, only reacts to variance of the contract in their own type, the buyer should only vary outcomes with the winning supplier's type.[1] Consequently, the buyer prefers a first- over a second-price auction when conducting a fixed-quantity or fixed-budget auction.[2]

2.3 Optimal Procurement

For ease of exposition, this section first discusses the optimal mechanism for quasi-linear preferences (Section 2.3.1), a quantity goal (Section 2.3.2), and a spending

[1]Esö and Futó (1999) show that the buyer would prefer to reduce uncertainty even further by making outcomes independent of the type of the winning supplier. This would require paying every potential supplier a transfer upon entry and, therefore, violate the assumption that the buyer only pays the winning supplier. This chapter maintains the assumption (PtW) as it seems to align better with real-world practices.

[2]The weak concavity of the utility function (2.1) in t stems from the penalty function regarding the spending goal. As one explanation for this functional form is a loss-averse buyer, a loss-averse auctioneer could be an additional explanation for the scarcity of second-price quantity auctions in the real world (see, Rothkopf *et al.*, 1990, for a discussion). However, weak concavity is not necessarily a consequence of loss aversion. For example, the S-shaped value function originally proposed by Kahneman and Tversky (1979) is not weakly concave. The weak concavity of reference-dependent utility persists if loss aversion is modeled through a kink in a piecewise linear reference-dependent utility component, as in much of the literature (e.g., Abeler *et al.*, 2011; Balzer *et al.*, 2022; Barberis *et al.*, 2001; Heidhues and Kőszegi, 2008; Marzilli Ericson and Fuster, 2011).

goal (Section 2.3.3) individually, before combining them to the general optimal mechanism in Section 2.3.4.

2.3.1 Optimal Procurement with Quasi-Linear Preferences

Against the backdrop of Remark 2.1, an optimal (direct revelation) mechanism solves

$$\max_{x(\theta),t(\theta)} \quad \int_{\underline{\theta}}^{\bar{\theta}} u(x(\theta)\,,t(\theta))\; f_1(\theta)\; d\theta \tag{2.2a}$$

$$\text{s.t.} \quad \pi(\theta,\tilde{\theta}=\theta) \geq \pi(\theta,\tilde{\theta}\neq\theta) \qquad \forall\theta,\tilde{\theta}\in[\underline{\theta};\bar{\theta}] \tag{2.2b}$$

$$\pi(\theta,\tilde{\theta}=\theta) \geq 0 \qquad \forall\theta \quad \in[\underline{\theta};\bar{\theta}], \tag{2.2c}$$

where $\pi(\theta,\tilde{\theta}) = (t(\tilde{\theta}) - \theta x(\tilde{\theta}))(1-F(\tilde{\theta}))^{n-1}$ is the expected profit of a supplier of type θ who announces to be of type $\tilde{\theta}$. The incentive compatibility constraint (2.2b) ensures that truthful reporting of types by suppliers constitutes an equilibrium: If all other suppliers report their true type, a supplier who reports being of type $\tilde{\theta}$ is awarded the contract with probability $(1-F(\tilde{\theta}))^{n-1}$ by assumption. Thus, (2.2b) guarantees that no bidder has an incentive to unilaterally deviate from truthful reporting. The participation constraint (2.2c) ensures that a supplier of type θ participates in the tender.

By standard steps, the incentive compatibility constraint can be re-expressed in terms of global and local incentive compatibility. Global incentive compatibility requires $\frac{d}{d\theta}\, x(\theta)\left(1-F(\theta)\right)^{n-1} \leq 0$. By requiring that the expected quantity sourced from a supplier of type θ is decreasing in θ, it ensures that inefficient types do not mimic more efficient types because this would require them to supply a higher *expected* quantity, which is disproportionately expensive for them. However, the buyer can implement a flat or even increasing schedule $x(\cdot)$ of *actual* procurement without violating global incentive compatibility because less efficient types are less likely to win the contract. Local incentive compatibility implies that

$$\pi(\theta) = \pi(\bar{\theta}) + \int_{\theta}^{\bar{\theta}} x(z) \left(1 - F(z)\right)^{n-1} dz. \tag{2.3}$$

Intuitively, a supplier of type θ can realise expected profits by pretending to be less efficient to the degree that they have lower production costs for the expected quantity demanded from a marginally less efficient type. A truthful mechanism must grant them the same level of expected profits as (expected) information rent. Thus, the quantity purchased from relatively inefficient types impacts the incentive-compatible payment to more efficient types but not vice versa. Equation (2.3) also implies that the least efficient type obtains the lowest expected profits. As a higher expected profit for suppliers implies a lower procured quantity or a higher transfer, i.e., a lower utility for the buyer, the participation constraint is binding for the least efficient type, and $\pi(\bar{\theta}) = 0$. Thus, the optimal transfer is

$$t^*(\theta) = \underbrace{\theta\, x(\theta)}_{\substack{\text{Production}\\\text{costs}}} + \underbrace{\int_{\theta}^{\bar{\theta}} x(z)\left(1 - F(z)\right)^{n-1} dz}_{\text{Expected information rent}} \underbrace{\left(1 - F(\theta)\right)^{1-n}}_{\substack{\text{Scaling to}\\\text{actual payment}}}. \tag{2.4}$$

The optimal transfer reimburses a supplier for their production costs and adds an incentive payment. The *actual* incentive payment to a supplier scales their expected information rent by their probability of winning the contract. As more efficient types obtain higher information rents but are also more likely to win, they do not necessarily obtain higher incentive payments. Hence, they only receive a higher transfer if they have to produce sufficiently many more units, which drives up their production costs, i.e., if the slope of the quantity schedule is sufficiently negative. If all types produce the same quantity, more efficient types obtain a lower transfer.

With the optimal transfer schedule from (2.4), the buyer's problem becomes

$$\begin{aligned} \max_{x(\theta)} \quad & \underbrace{\int_{\underline{\theta}}^{\bar{\theta}} \left(v\left(x(\theta)\right) - \left(\theta + \frac{F(\theta)}{f(\theta)}\right) x(\theta) \right) f_1(\theta) d\theta}_{\text{Direct utility}} \\ & \quad \underbrace{- \int_{\underline{\theta}}^{\bar{\theta}} \lambda_Q \mathbf{1}_Q(\cdot) \left(\hat{x} - x(\theta)\right) f_1(\theta) d\theta}_{\text{Penalty for missing quantity goal}} \\ & \qquad \underbrace{- \int_{\underline{\theta}}^{\bar{\theta}} \lambda_B \left(\left(\mathbf{1}_B \, \theta + \frac{\int_{\underline{\theta}}^{\theta} \mathbf{1}_B(\cdot) \, f(z) \, dz}{f(\theta)} \right) x(\theta) - \hat{t} \right) f_1(\theta) d\theta}_{\text{Penalty for missing spending goal}} \end{aligned} \tag{2.5a}$$

$$\text{s.t.} \qquad \frac{d}{d\theta} \, x(\theta) \left(1 - F(\theta)\right)^{n-1} \leq 0. \tag{2.5b}$$

The third line in (2.5a) reflects the potential disutility from missing the spending goal. It shows that there are two ways, in which procuring more from a supplier of type θ impacts the penalty for overspending: First, it raises the buyer's spending on procurement from that type by θ, which corresponds to the supplier's additional production costs. This causes the buyer to suffer from a penalty for overspending if and only if they are spending more than $\hat{t}$ on procurement *from that type*. Second, due to the structure of the incentive-compatible transfer from (2.4), procuring more from type θ requires paying higher information rents to more efficient suppliers. This results in a penalty if and only if the buyer spends more than $\hat{t}$ from *more efficient* types than θ.

For the case of quasi-linear preferences, we have $\mathbf{1}_Q = \mathbf{1}_B = 0$, such that the second and third line drop out of the expression in (2.5a). However, to avoid repetition later, consider the slightly more general case of constant $\mathbf{1}_Q$ and $\mathbf{1}_B$. If $\mathbf{1}_B$ is constant, the third line in (2.5a) resolves to $-\int_{\underline{\theta}}^{\bar{\theta}} \lambda_B \, \mathbf{1}_B \left(\theta + \frac{F(\theta)}{f(\theta)}\right) x(\theta) f_1(\theta) d\theta$. The optimal mechanism emerges from a point-wise maximisation of the integrand in (2.5a). Consequently, the optimal "baseline" quantity schedule (subscript BL) for constant $\mathbf{1}_Q$ and $\mathbf{1}_B$, $x^*_{BL}(\theta)$, solves

$$v'(x^*_{BL}(\theta)) = (1 + \lambda_B \, \mathbf{1}_B) \left(\theta + \frac{F(\theta)}{f(\theta)}\right) - \lambda_Q \, \mathbf{1}_Q. \tag{2.6}$$

In the case of quasi-linear preferences (or irrelevant goals), the optimal quantity schedule corresponds to (2.6) with $\mathbf{1}_B = \mathbf{1}_Q = 0$. If the indicator variables are constant, the monotone reverse hazard rate property ensures that the quantity schedule in (2.6) is decreasing and, therefore, satisfies the global incentive compatibility constraint from (2.5b).

2.3.2 Optimal Procurement with a Quantity Goal

The effects of adding a quantity goal depend on how demanding it is.

Optimal Mechanism with a Very Lenient or Very Demanding Quantity Goal.— Equation (2.6) already covers the cases of a very lenient and a very demanding quantity goal. If the quantity goal is very lenient, the buyer meets the target for every type of the winning supplier. Thus, the optimal quantity schedule corresponds to (2.6) with $\mathbf{1}_Q = 0$, which in the following is denoted by $x^*_{BL;\mathbf{1}_Q=0}(\theta)$. Similarly, if the goal is very demanding, the buyer never meets the target and the optimal quantity schedule corresponds to $x^*_{BL;\mathbf{1}_Q=1}(\theta)$. Therefore, the optimal quantity schedule only differs from the one for quasi-linear preferences if the buyer globally misses their quantity goal. In this case, the buyer increases the procured quantity. Intuitively, when they miss their quantity goal, the buyer has a higher marginal utility of raising the procured quantity because this additionally results in a lower quantity-related penalty. This results in a discrete upward shift of the quantity schedule.

Very Demanding Quantity Goal: The Limiting Case of Fixed-Quantity Procurement.— Consider the marginal utility of procuring an additional unit from a given type θ. Ignoring the probability weight, it is

$$\begin{cases} v'(x) + \lambda_Q - (1+\lambda_B\,\mathbf{1}_B)\left(\theta + \frac{F(\theta)}{f(\theta)}\right) & \text{for } x \leq \hat{x} \\ v'(x) \qquad\quad - (1+\lambda_B\,\mathbf{1}_B)\left(\theta + \frac{F(\theta)}{f(\theta)}\right) & \text{for } x > \hat{x}. \end{cases} \tag{2.7}$$

If the buyer procures less than their target quantity, additional procurement helps to reduce the penalty for overspending by λ_Q. This positive effect on the marginal utility disappears once they buy more than $\hat{x}$; the marginal utility jumps down at the target quantity.

Consequently, if the quantity goal is important enough, i.e., λ_Q is high enough, it can be that the second line in (2.7) is negative when evaluated at $x = \hat{x}$ for all types, while, at the same time, the first line in (2.7) is positive for all θ when evaluated at $\hat{x}$. Additional procurement is profitable if and only if it goes along with a reduction in the penalty for missing the quantity goal. Thus, the buyer has a positive marginal utility of procuring more if $x \leq \hat{x}$ and a negative one if $x > \hat{x}$ for all types of the winning supplier. Consequently, the buyer procures their target quantity from all types. The corresponding optimal transfer to a winning supplier of a given type is identical to the bid of that type in a first-price fixed-quantity auction without a reserve price. Thus, the optimality of fixed-quantity auctions is a special case of this model.

Excluding Fixed-Quantity Procurement.— Fixed-quantity procurement is the result of a demanding quantity goal of high importance. First, the buyer must miss the goal when ignoring it, i.e., the quantity goal $\hat{x}$ must be sufficiently high such that the marginal utility of buying the $\hat{x}$-th unit in the second line of (2.7) is negative for all types. Second, the quantity goal must be so important that the buyer procures at least $\hat{x}$ units when they account for the goal, i.e., the marginal utility of buying the $\hat{x}$-th unit in the first line of (2.7) must be positive for all types. As the primary focus of this chapter are explicit goals outside the well-known special cases, the rest of this chapter assumes that the latter does not hold, i.e.,

$$\lambda_Q < (1 + \lambda_B \, \mathbf{1}_B) \left(\bar{\theta} + \frac{1}{f(\bar{\theta})} \right) - v'(\hat{x}). \tag{2.8}$$

Assumption (2.8) implies that the buyer procures less than $\hat{x}$ at least from the least efficient type; intuitively, the quantity goal is sufficiently unimportant for the buyer to accept buying less than $\hat{x}$ at least from the type from which they want to source the lowest quantity. Hence, (2.8) rules out not only fixed-quantity procurement but also the special case of quasi-linear procurement because it ensures that the quantity goal is actually relevant.

Moderately Demanding Quantity Goal: Partial Bunching.— By assumption (2.8), the buyer sometimes falls short of their target quantity. This assumption does not imply that the goal is so demanding that they always do. So far, the analysis has neglected cases, where the buyer misses their quantity goal only when sourcing from some types. Given that $x^*_{BL}(\theta)$ is decreasing in θ, the buyer will miss their quantity goal only when procuring from inefficient suppliers, from whom the buyer does not source as much. Formally, consider the case where $x^*_{BL;\mathbf{1}_Q=0}(\theta_1) = \hat{x}$ and $x^*_{BL;\mathbf{1}_Q=1}(\theta_2) = \hat{x}$ for $\theta_2 \in (\theta_1, \bar{\theta})$. As illustrated in the left panel of Figure 2.2, not accounting for the penalty for missing the goal – i.e., following the quantity schedule $x^*_{BL;\mathbf{1}_Q=0}(\theta)$ – indeed prescribes a quantity above $\hat{x}$ for $\theta \leq \theta_1$. For $\theta \geq \theta_2$, the buyer sources less than $\hat{x}$, even when accounting for the associated disutility, i.e., $x^*_{BL;\mathbf{1}_Q=1}(\theta) \leq \hat{x}$. The interval (θ_1, θ_2) corresponds to the case of fixed-quantity procurement: The buyer has a positive marginal utility of procuring more if $x \leq \hat{x}$ and a negative one if $x > \hat{x}$. Thus, the optimal quantity schedule with a quantity goal (subscript Q), $x^*_Q(\theta)$, emerges as

$$x^*_Q(\theta) = \begin{cases} x^*_{BL;\mathbf{1}_Q=0}(\theta) & \text{for } \theta \leq \theta_1 \\ \hat{x} & \text{for } \theta \in (\theta_1, \theta_2) \\ x^*_{BL;\mathbf{1}_Q=1}(\theta) & \text{for } \theta \geq \theta_2. \end{cases} \tag{2.9}$$

Figure 2.2 illustrates the quantity schedule $x^*_Q(\theta)$ and the corresponding transfers.

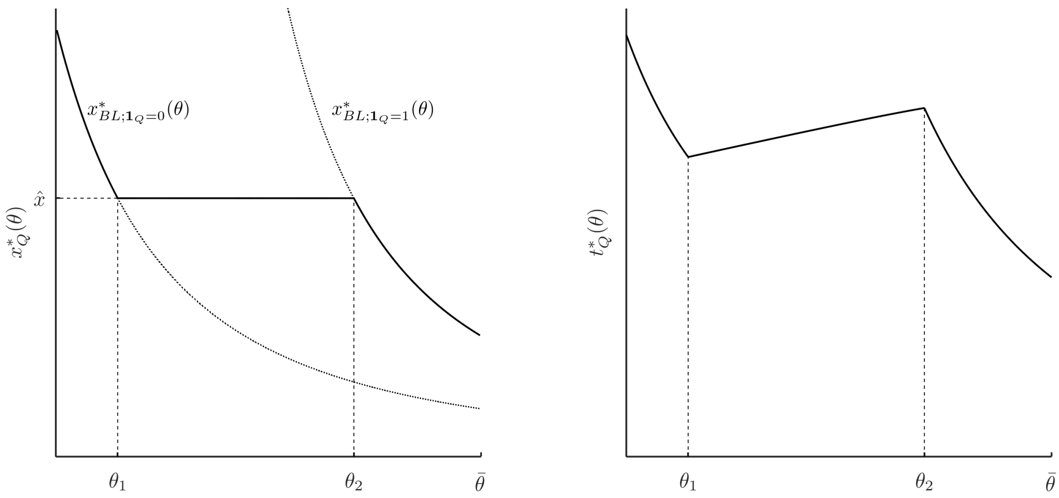

Figure 2.2: Optimal procurement mechanisms $(x^*_Q(\theta), t^*_Q(\theta))$ for a buyer with a relevant but non-dominant quantity goal.

2.3.3 Optimal Procurement with a Spending Goal

Next, consider the effect of a spending goal.

Optimal Mechanism with a Very Lenient or Very Demanding Spending Goal.— As with a quantity goal, the optimal quantity schedule with a spending goal corresponds to the schedule from (2.6), as long as $\mathbf{1}_B$ is constant. If the buyer always spends less than envisaged, we have $\mathbf{1}_B = 0$. A buyer whose spending always exceeds their spending goal follows the schedule from (2.6) with $\mathbf{1}_B = 1$. If the buyer globally misses their spending goal, this results in a discrete downward shift of the procured quantity, i.e., $x^*_{BL;\mathbf{1}_B=1}(\theta) < x^*_{BL;\mathbf{1}_B=0}(\theta)$. Higher procured quantities result in higher spending, such that a buyer suffering from overspending curbs the procured quantity.

Very Demanding Spending Goal: The Limiting Case of Fixed-Budget Procurement.— As with a quantity goal, the optimal mechanism collapses into fixed-budget procurement if the penalty for missing the spending goal is sufficiently large and the spending goal sufficiently demanding. The marginal utility of procuring an additional unit (ignoring its probability weights) is

$$\begin{cases} v'(x) - \left(\theta + \frac{F(\theta)}{f(\theta)}\right) + \lambda_Q \, \mathbf{1}_Q - \frac{\int_{\underline{\theta}}^{\theta} \lambda_B \mathbf{1}_B(\cdot) \, f(z) \, dz}{f(\theta)} & \text{for } x < x_{FB}(\theta) \\ v'(x) - \left(\theta + \frac{F(\theta)}{f(\theta)}\right) + \lambda_Q \, \mathbf{1}_Q - \frac{\int_{\underline{\theta}}^{\theta} \lambda_B \, \mathbf{1}_B(\cdot) \, f(z) \, dz}{f(\theta)} - \lambda_B \, \theta & \text{for } x \geq x_{FB}(\theta) \end{cases} \tag{2.10}$$

where $x_{FB}(\theta)$ (subscript FB for "fixed budget") is the quantity for which spending is equal to $\hat{t}$ given the quantity schedule in the interval $[\theta, \bar{\theta}]$. The buyer always accounts for the effect of additional procurement on their direct utility and the potential effect on the penalty for missing their quantity target. Furthermore, buying more has the negative effect of increasing the information rents of more efficient types, which increases the (potential) penalty for overspending when sourcing from them. However, there is a jump in the marginal utility at the target spending. If the buyer overspends, i.e., $x \geq x_{FB}(\theta)$, paying the additional production costs of θ comes with a higher disutility, as it increases the penalty for overspending, which has a weight of λ_B.

This jump in the marginal utility results in bunching. It can be that the buyer has a positive marginal utility of buying more if they spend less than $\hat{t}$ but a negative marginal utility when they overspend. As long as this is the case, they spend exactly their budget. Thus, the size of the bunching interval increases with the size of the jump in the marginal utility, i.e., with λ_B. With a high enough importance of the spending goal, λ_B, it can even be that the first line in (2.10) is positive for all types, whereas the second line in (2.10) is always negative. In such a case, the buyer should engage in fixed-budget procurement by always spending $\hat{t}$. Generally, to keep spending at the level t, the quantity schedule must satisfy the differential equation $x'_{FB}(\theta) = (x_{FB}(\theta) - t/\theta)(n-1)f(\theta)/(1-F(\theta))$. To globally have spending of $\hat{t}$, the buyer sources

$$x^*_{FB}(\theta) = \frac{C - \int_{\underline{\theta}}^{\theta} \frac{\hat{t}}{z} \, (n-1) \left(1 - F(z)\right)^{n-2} f(z) \, dz}{\left(1 - F(\theta)\right)^{n-1}} \tag{2.11}$$

units, where the arbitrary integration constant C is such that $x^*_{FB}(\bar{\theta}) = \hat{t}/\bar{\theta}$: The buyer spends $\hat{t}$ on procurement from type $\bar{\theta}$ and this type makes zero profits.[3] For the limiting case of fixed-budget procurement, the optimal mechanism corresponds to a first-price budget auction without a reserve price: The optimal quantity schedule from (2.11) corresponds to the quantity a supplier of type θ would bid in such an auction.

Excluding Fixed-Budget Procurement.— Fixed-budget procurement requires that target spending, and, therefore, the "admissible" quantity $x_{FB}(\theta)$ is low enough for the buyer to procure more than $x_{FB}(\theta)$ units when ignoring the goal. Formally, $x_{FB}(\theta)$ must be so low that the first line in (2.10) is positive for all types when evaluated at $x_{FB}(\theta)$. Furthermore, the budget goal must be so important that the buyer does not overspend, when they account for their spending goal. Formally, λ_B must be high enough for the marginal utility in the second line of (2.10) to be negative for the $x_{FB}(\theta)$-th unit. Again, to focus on the novel aspects, the rest of this chapter excludes fixed-budget procurement by assuming that the buyer does not attach such a high value to their goal. Instead, the buyer accepts overspending on some type. As discussed in Section 2.3.1, this can be any type: There is no guarantee that the buyer sources *sufficiently much* more from efficient types to spend more when contracting with them. However, as we will see, the analysis of the effects of a spending goal is particularly interesting when the buyer pays a higher transfer to more efficient types under the baseline mechanism. Formally, this requires that the buyer does not suffer too much from variations in the procured

[3]Interestingly, fixed-budget procurement results in "no distortion at the bottom:" Buying $\hat{t}/\bar{\theta}$ units from the type $\bar{\theta}$ corresponds to the first-best quantity, in the sense that it maximises the buyer's objective function subject to the suppliers' participation constraint.

quantity, i.e., that

$$-v''(x) < \frac{(1+\lambda_B\,\mathbf{1}_B)\left(1+\frac{d}{d\theta}\frac{F(\theta)}{f(\theta)}\right)}{\left(\hat{t}/\theta - x\right)\frac{(n-1)f(\theta)}{(1-F(\theta))}} \text{ for } x \in \{x^*_{BL;\mathbf{1}_B=1}(\theta); x^*_{BL;\mathbf{1}_B=0}(\theta)\}.^4 \quad (2.12)$$

Under assumption (2.12), analysing spending goals that are only partially relevant, requires that the buyer overspends at least when trading with the most efficient type. Therefore, the rest of this chapter assumes that

$$\lambda_B < \frac{v'\left(\hat{t}/\underline{\theta}\right) + \lambda_Q\,\mathbf{1}_Q - \underline{\theta}}{\underline{\theta}}.^5 \quad (2.13)$$

Moderately Demanding Spending Goal: Partial Bunching.— Consequently, when the spending goal is moderately demanding such that the buyer only overspends when sourcing from some supplier types, they will overspend when sourcing from efficient types but not when sourcing from inefficient types. Because the marginal utility of additional procurement jumps up when spending reaches $\hat{t}$, there is a bunching interval in between: Once spending reaches $\hat{t}$ for an efficient enough type, spending marginally more has a negative marginal utility as it now results in a penalty on the additional production costs. Hence, the optimal quantity schedule with a budget goal (subscript B) takes the form

$$x^*_B(\theta) = \begin{cases} x^*_{BL;\mathbf{1}_B=1}(\theta) & \text{for } \theta \le \theta_3 \\ x^*_{FB}(\theta) \text{ with } x^*_{FB}(\theta_3) = x^*_{BL;\mathbf{1}_B=1}(\theta_3) & \text{for } \theta \in (\theta_3, \theta_4) \\ x^*_{BB}(\theta) & \text{for } \theta \ge \theta_4, \end{cases} \quad (2.14)$$

where spending weakly exceeds $\hat{t}$ for $\theta \le \theta_3$ and is weakly below $\hat{t}$ for $\theta \ge \theta_4$.

First, consider the shape of the quantity schedules for given interval boundaries. For $\theta \le \theta_3$, the marginal effects of raising the procured quantity are identical to the case where $\mathbf{1}_B = 1$ for all types. Intuitively, the buyer suffers from the penalty

[4] A decreasing transfer requires that $|x'_{BL}(\theta)| > |x'_{FB}(\theta)|$, where $x'_{BL}(\theta)$ can be obtained from differentiating the first-order condition (2.6).

[5] Assumption (2.13) is a sufficient condition in the sense that we must have $x_{BL}(\underline{\theta}) > x_{FB}(\underline{\theta})$ and given that type $\underline{\theta}$ obtains a positive information rent we have $x_{FB}(\underline{\theta}) < \hat{t}/\underline{\theta}$.

for overspending when procuring more and misses their spending goal whenever sourcing from a more efficient type in either case; formally, the second line in the marginal utility from (2.10) resolves to $v'(x) - (1 + \lambda_B)(\theta + F(\theta)/f(\theta)) + \lambda_Q \mathbf{1}_Q$. The buyer overspends if, even when accounting for those negative effects of higher procurement, setting this marginal utility to zero results in overspending. Hence, it must be that $x^*_{FB}(\theta_3) = x^*_{BL;\mathbf{1}_B=1}(\theta_3)$. Taking the interval boundary θ_3 as given, this condition pins down the integration constant in the quantity schedule $x^*_{FB}(\theta)$.

The "below budget"(subscript BB) quantity schedule $x^*_{BB}(\theta)$ maximises those parts of the buyer's expected utility from (2.5a) that depend on the quantity schedule in the interval $[\theta_4, \bar{\theta}]$. While additional procurement does not result in a penalty on the additional production costs, the buyer must account for the fact that the additional information rents increase the penalty for overspending when they source from types $\theta \leq \theta_3$. Formally, the expected penalty for overspending in the third line of (2.5a) resolves to

$$-\int_{\underline{\theta}}^{\theta_3} \lambda_B \left(\left(\theta + \frac{F(\theta)}{f(\theta)} \right) x(\theta) - \hat{t} \right) f_1(\theta) d\theta - \int_{\theta_3}^{\bar{\theta}} \lambda_B \left(\frac{F(\theta_3)}{f(\theta)} x(\theta) - \hat{t} \right) f_1(\theta) d\theta$$

and, therefore, depends on the quantity schedule for $\theta \geq \theta_4$. Given the structure of the optimal mechanism from (2.14), the buyer switches to constant spending at the end of the below-budget interval. Thus, their maximization problem for determining the optimal quantity schedule for $\theta \geq \theta_4$ comes with two boundary conditions. First, as seen before, keeping spending constant at $\hat{t}$ requires following a continuous quantity schedule described by a differential equation. Thus, it must be that $x^*_{BB}(\theta_4) = x^*_{FB}(\theta_4)$. Second, spending must reach $\hat{t}$ under the quantity schedule $x^*_{BB}(\theta)$ in the first place, i.e., $t^*(x^*_{BB}(\theta_4)) = \hat{t}$. Hence, the optimal choice

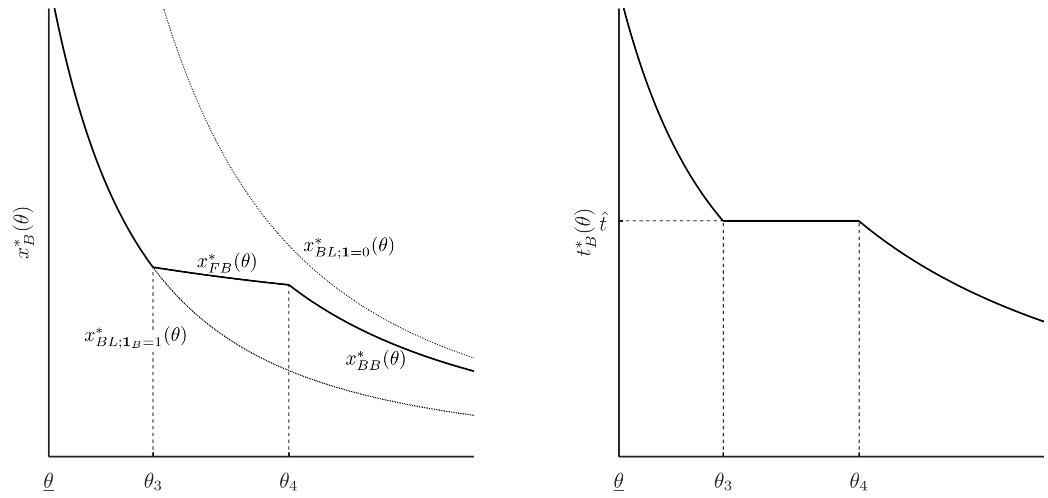

Figure 2.3: Optimal procurement mechanisms $(x^*_B(\theta), t^*_B(\theta))$ for a buyer with a relevant but non-dominant spending goal.

of $x^*_{BB}(\theta)$ amounts to the isoperimetric variational calculus problem

$$\max_{x_{BB}(\theta)} \int_{\theta_4}^{\bar{\theta}} \underbrace{\left\{ v\left(x_{BB}(\theta)\right) + \left(\lambda_Q \,\mathbf{1}_Q - \left(\theta + \frac{F(\theta)}{f(\theta)}\right)\right) x_{BB}(\theta) \right\} f_1(\theta)}_{\text{Utility in } [\theta_4, \bar{\theta}]}$$

$$\underbrace{- \frac{\lambda_B \, F(\theta_3)}{f(\theta)} \, x_{BB}(\theta) \, f_1(\theta)}_{\text{Effect on utility in } [\underline{\theta}, \theta_3]} \; d\theta$$

$$\text{s.t.} \qquad \int_{\theta_4}^{\bar{\theta}} x_{BB}(\theta) \, \Big(1 - F(\theta)\Big)^{n-1} \, d\theta = \Big(\hat{t} - \theta_4 \, x_{FB}(\theta_4)\Big) \Big(1 - F(\theta_4)\Big)^{n-1}.$$

The optimising path is

$$v'\left(x^*_{BB}(\theta)\right) = \left(\theta + \frac{F(\theta)}{f(\theta)}\right) - \lambda_Q \mathbf{1}_Q + \frac{\lambda/n + \lambda_B \, F(\theta_3)}{f(\theta)}, \tag{2.15}$$

where λ corresponds to the Lagrange multiplier of the constraint on total spending. As Figure 2.3 illustrates, the resulting quantity schedule procures less than a buyer without a relevant spending goal but more than one who was already overspending would, i.e., $x^*_{BL;\mathbf{1}_B=1}(\theta) < x^*_{BB}(\theta) < x^*_{BL;\mathbf{1}_B=0}(\theta)$. Formally, this is because the shadow price keeps within certain bounds, as Appendix 2.A shows. Intuitively, the buyer buys less than they would if they were already overspending because the "current" production costs are not subject to a penalty for overspending. However, they curb their demand relative to a situation where the spending goal is irrelevant to lower the penalty associated with higher information rents for more efficient suppliers. This contrasts with quantity goals, where not missing the goal is equivalent to not having a relevant goal.

The procured quantity in the below-budget interval only impacts the penalty for overspending if the buyer spends more on procurement from more efficient types. Thus, the optimal mechanism simplifies if the buyer overspends only when procuring from relatively inefficient types: When spending less than their budget the buyer does not change the information rents of the less efficient types on which they overspend. Consequently, the quantity schedule when spending less than $\hat{t}$ is $x^*_{BL;\mathbf{1}_B=0}(\theta)$. Hence, considering these cases does not add much insight, which is why this chapter assumes that assumption (2.12) holds and the transfer is decreasing under the baseline mechanism.

Due to the additional term on the right-hand side of equation (2.15), the monotone reverse hazard rate property, is not sufficient to ensure that the below-budget quantity schedule satisfies global incentive compatibility. Consequently, this chapter assumes that the sharpened monotone reverse hazard rate property

$$\frac{d}{d\theta}\frac{F(\theta)}{f(\theta)} > \max\left\{\lambda_B \frac{\big(\theta_4\, f(\theta_4) + F(\theta_4)\big) f'(\theta)}{(f(\theta))^2} - 1; 0\right\} \tag{2.16}$$

holds.[6] Just like the standard monotone hazard rate property, condition (2.16) implies that the right-hand side of (2.15) is increasing in θ such that the quantity schedule $x^*_{BB}(\theta)$ is decreasing in θ.

[6]This conditions ensures an increasing right-hand side of (2.15) even when λ takes its highest possible value. Sufficient conditions in terms of model fundamentals are, for example, $f'(\theta) < 0$, in which case (2.16) resolves to the standard monotone reverse hazard property, or $\frac{d}{d\theta}\frac{F(\theta)}{f(\theta)} > \max\{\lambda_B\big(\theta\, f(\theta) + F(\theta)\big) f'(\theta)/(f(\theta))^2 - 1; 0\}$ if $\theta f(\theta) + F(\theta)$ is increasing in θ.

So far, we have taken θ_3 as given. In the optimal mechanism, the buyer chooses θ_3 to maximise their expected utility. The other parameters are pinned down by the structure of the optimal mechanism in (2.14). For the buyer to spend weakly more than $\hat{t}$ for $\theta \leq \theta_3$, even when accounting for the penalty for overspending, C must satisfy

$$x^*_{FB}(\theta_3; C) = x^*_{BL;\mathbf{1}_B=1}(\theta_3). \tag{2.17a}$$

Keeping spending constant at $\hat{t}$ requires following a continuous quantity schedule. Hence, θ_4 is determined by

$$x^*_{FB}(\theta_4; C) = x^*_{BB}(\theta_4; \lambda, \theta_3). \tag{2.17b}$$

Finally, by construction, λ must be such that $t(\theta_4) = \hat{t}$, i.e.,

$$\hat{t} = \theta_4 x^*_{BB}(\theta_4; \lambda, \theta_3) + \int_{\theta_4}^{\bar{\theta}} x^*_{BB}(\theta; \lambda, \theta_3) \left(1 - F(\theta)\right)^{n-1} d\theta \left(1 - F(\theta_4)\right)^{1-n}. \tag{2.17c}$$

To determine the optimal mechanism, the buyer solves $\frac{d\mathbb{E}[U]}{d\theta_3} = 0$ with

$$\begin{aligned}\frac{d\mathbb{E}[U]}{d\theta_3} = & \int_{\theta_3}^{\theta_4} \frac{\partial \Psi}{\partial x^*_{FB}} \frac{\partial x^*_{FB}}{\partial C} \frac{dC}{d\theta_3} f_1(\theta)\, d\theta + \int_{\theta_4}^{\bar{\theta}} \frac{\partial \Psi}{\partial x^*_{BB}} \left(\frac{\partial x^*_{BB}}{\partial \theta_3} + \frac{\partial x^*_{BB}}{\partial \lambda} \frac{d\lambda}{d\theta_3} \right) f_1(\theta)\, d\theta \\ & - \lambda_B \int_{\theta_3}^{\theta_4} \left(\frac{F(\theta_3)}{f(\theta)} \right) \frac{\partial x^*_{FB}}{\partial C} \frac{dC}{d\theta_3} f_1(\theta)\, d\theta \\ & - \lambda_B \int_{\theta_4}^{\bar{\theta}} \left(\frac{F(\theta_3)}{f(\theta)} \right) \left(\frac{\partial x^*_{BB}}{\partial \theta_3} + \frac{\partial x^*_{BB}}{\partial \lambda} \frac{d\lambda}{d\theta_3} \right) f_1(\theta)\, d\theta \\ & + \lambda_B \left(\frac{F(\theta_3)}{f(\theta_3)} \right) x^*_{FB}(\theta_3)\, f_1(\theta_3),\end{aligned}$$

where $\Psi = v\left(x^*_{FB/BB}(\theta)\right) + \left(\lambda_Q \mathbf{1}_Q - \left(\theta + \frac{F(\theta)}{f(\theta)}\right)\right) x^*_{FB/BB}(\theta)$. The relevant derivatives can be found in equations (2.A.1) – (2.A.8) in Appendix 2.A, where those regarding C, θ_4, and λ result from implicit differentiation of the system in (2.17).

2.3.4 Optimal Procurement with a Quantity and a Spending Goal

The optimal mechanism when the buyer has both a quantity and a spending goal results from combining the insights from the previous sections. By assumptions (2.8), (2.12), and (2.13), both goals are relevant; for both goals, there are at least some types for which the buyer misses the goal. Consequently, there are four possible cases:

Very Demanding Quantity and Spending Goal.— If both goals are very demanding, the buyer procures too little and spends too much sourcing from any type. Thus, $\mathbf{1}_Q = \mathbf{1}_B = 1$, such that the optimal quantity schedule corresponds to $x^*_{BL;\mathbf{1}_Q=1,\mathbf{1}_B=1}(\theta)$.

Moderately Demanding Quantity Goal and Very Demanding Spending Goal.— If the quantity goal is somewhat less demanding such that the buyer sometimes meets their quantity target, the basic structure of the optimal quantity schedule must follow (2.9). Coupled with globally missing the spending goal, the optimal quantity schedule is $x^*_{Q;\mathbf{1}_B=1}(\theta)$.

Very Demanding Quantity Goal and Moderately Demanding Spending Goal.— Analogously, the structure of the optimal quantity schedule when the buyer sometimes overspends is (2.14); if they always miss their quantity goal, the optimal quantity schedule is $x^*_{B;\mathbf{1}_Q=1}(\theta)$.

Moderately Demanding Quantity and Spending Goal.— When both goals are moderately demanding, such that the buyer does not miss either goal globally, two new cases emerge. Under assumptions (2.8) and (2.13), we know that the buyer spends weakly less than $\hat{t}$ and procures less than $\hat{x}$ from the least efficient types, whereas they spend more than $\hat{t}$, buying weakly more than $\hat{x}$ from the most efficient types. As seen in Sections 2.3.2 and 2.3.3, the buyer engages in partial bunching in the respective dimensions, whenever they switch between missing and not missing the goal. Thus, the main question, when the buyer misses either goal sometimes, is the order of the bunching intervals. First, suppose that the quantity bunching interval occurs when the buyer overspends. Figure 2.4 illustrates the optimal mechanism for this case. The fact that the quantity schedule is partially flat in the interval where the buyer spends more than $\hat{t}$ does not change the fundamental logic of the quantity schedule from Section 2.3.3: The marginal penalty associated with overspending is unchanged. The way in which the quantity schedule when the buyer does not miss their spending goal impacts total expenditures when they do overspend is unchanged, as well. Similarly, the effect of the quantity goal is just as in Section 2.3.2. When the buyer switches between missing and not missing their quantity goal, while also missing their spending goal, this results in partial bunching. Thus, the optimal quantity schedule is

$$x^*_{QB}(\theta) = \begin{cases} x^*_{BL;\mathbf{1}_B=1,\mathbf{1}_Q=0}(\theta) & \text{for } \theta \leq \theta_1 \\ \hat{x} & \text{for } \theta \in (\theta_1, \theta_2) \\ x^*_{BL;\mathbf{1}_B=1,\mathbf{1}_Q=1}(\theta) & \text{for } \theta \in [\theta_2, \theta_3] \\ x^*_{FB}(\theta) \text{ with } x^*_{FB}(\theta_3) = x^*_{BL;\mathbf{1}_B=1,\mathbf{1}_Q=1}(\theta_3) & \text{for } \theta \in (\theta_3, \theta_4) \\ x^*_{BB;\mathbf{1}_Q=1}(\theta) & \text{for } \theta \geq \theta_4, \end{cases} \tag{2.18}$$

where $x^*_{BL;\mathbf{1}_B=1,\mathbf{1}_Q=0}(\theta_1) = x^*_{BL;\mathbf{1}_B=1,\mathbf{1}_Q=1}(\theta_2) = \hat{x}$.

The formulation in (2.18) assumes that the buyer spends more than $\hat{t}$ for $\theta < \theta_3$. This need not be the case: With a flat quantity schedule, more efficient types obtain lower transfers such that spending could drop below $\hat{t}$ in the quantity bunching interval. In that case, there is a type, for which $t(\theta) = \hat{t}$ and $x(\theta) = \hat{x}$, which requires deciding which goal dominates given that exactly hitting the goal usually results in bunching in the respective dimension. Something similar happens if the two bunching intervals coincide in the sense that there is a type for which

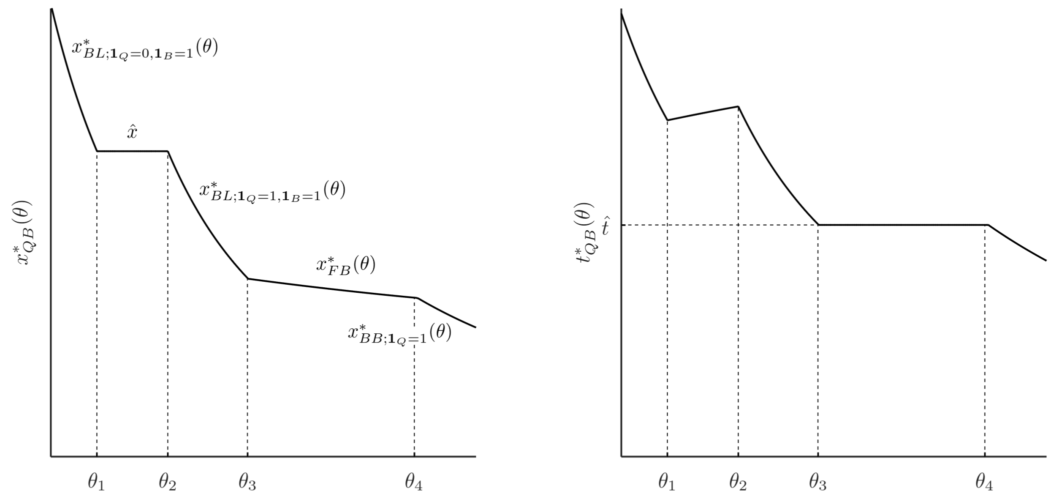

Figure 2.4: Optimal mechanism $(x^*_{QB}(\theta), t^*_{QB}(\theta))$ for a buyer with moderately demanding quantity and spending goals when the quantity bunching interval occurs while the buyer overspends.

$x^*_{FB}(\theta) = \hat{x}$ in the interval (θ_3, θ_4). As these cases do not produce additional insights, they are relegated to Appendix 2.A.

Instead, suppose that the quantity bunching interval occurs when expenditures are below $\hat{t}$. Figure 2.5 illustrates the optimal mechanism in this situation. The effects of the spending goal are unchanged to Section 2.3.3. When sourcing from inefficient types, the buyer spends less than $\hat{t}$ and accounts for the fact that increasing the procured quantity entails a higher penalty from overspending when they source from efficient types. The only difference is that they reach their quantity target under the quantity schedule $x^*_{BB}(\theta)$ from (2.15). The effect of the quantity goal is just as under the baseline quantity schedule. When the buyer stops falling short of their quantity target, the marginal utility of additional procurement jumps down. They engage in partial bunching regarding the quantity until buying more than $\hat{x}$ units is profitable even with the lower marginal utility. Hence, the

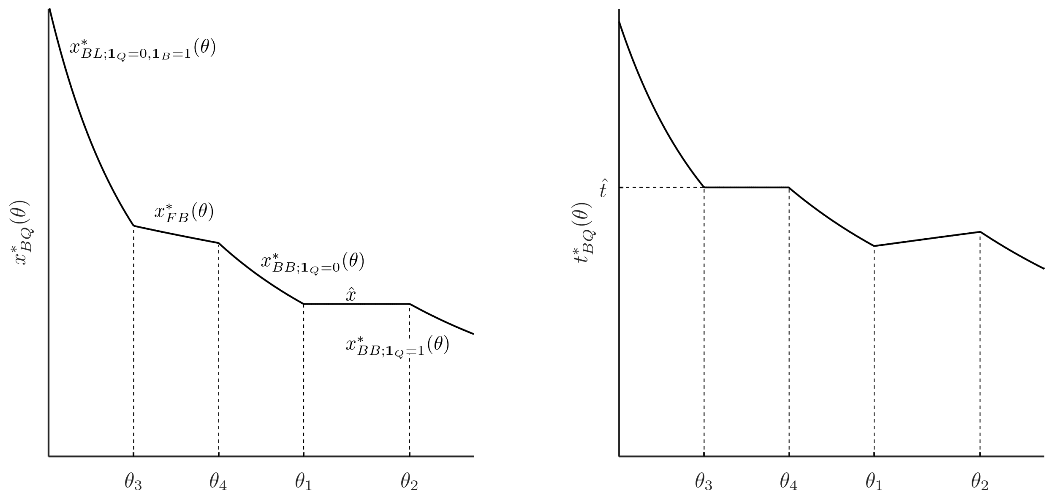

Figure 2.5: Optimal mechanism $(x^*_{BQ}(\theta), t^*_{BQ}(\theta))$ for a buyer with moderately demanding quantity and spending goals when the quantity bunching interval occurs while the buyer does not overspend.

optimal quantity schedule is

$$x^*_{BQ}(\theta) = \begin{cases} x^*_{BL;\mathbf{1}_B=1,\mathbf{1}_Q=0}(\theta) & \text{for } \theta \leq \theta_3 \\ x^*_{FB}(\theta) & \text{for } \theta \in (\theta_3, \theta_4) \\ x^*_{BB;\mathbf{1}_Q=0}(\theta) & \text{for } \theta \in [\theta_4, \theta_1] \\ \hat{x} & \text{for } \theta \in (\theta_1, \theta_2) \\ x^*_{BB;\mathbf{1}_Q=1}(\theta) & \text{for } \theta \geq \theta_2, \end{cases} \tag{2.19}$$

where, $x^*_{BB;\mathbf{1}_Q=1}(\theta_2) = \hat{x} = x^*_{BB;\mathbf{1}_Q=0}(\theta_1)$, $t^*(\theta_4) = \hat{t}$, $x_{FB}(\theta_4) = x_{BB}(\theta_4)$, and $x_{FB}(\theta_3) = x_{BL;\mathbf{1}_B=1,\mathbf{1}_Q=0}(\theta_3)$. Because more efficient types obtain lower transfers when the quantity schedule is flat, spending remains below $\hat{t}$ for $\theta \geq \theta_4$.

The preceding discussion has provided all essential insights into the structure and properties of the optimal procurement mechanism. Proposition 2.1 consolidates these insights into a comprehensive characterisation of the optimal mechanism.

Proposition 2.1 (Optimal Mechanism) *Consider an optimal mechanism for a buyer with utility function* (2.1) *under the constraints (LCW) and (PtW). Assume that the buyer sources from one of n suppliers, whose constant unit costs θ are* *private information and represent independent draws from a distribution* $F(\theta)$ *on* $[\underline{\theta}, \bar{\theta}]$ *with* $0 < \underline{\theta} < \bar{\theta}$*, where the sharpened monotone inverse hazard rate from* (2.16) *holds. Assume that conditions* (2.8) *and* (2.13) *are satisfied such that the buyer misses either goal at least sometimes. If the buyer's utility is not too concave in the sense of* (2.12)*, an optimal mechanism is characterised by a transfer schedule satisfying* (2.4) *and a quantity schedule that corresponds to one of the cases from* (2.6) *with* $\mathbf{1}_Q = \mathbf{1}_B = 1$*,* (2.9) *with* $\mathbf{1}_B = 1$*,* (2.14) *with* $\mathbf{1}_Q = 1$*,* (2.18)*,* (2.19)*, or* (2.A.9) – (2.A.11)*.*

2.4 Discussion

This chapter shows that considering explicit quantity and spending goals makes bunching, i.e., treating different supplier types the same regarding the procured quantity or the transfer paid to them, a central element of optimal procurement.

Traditionally, bunching (in terms of the quantity) is a way to deal with non-monotonic virtual costs. In that case, the global incentive compatibility becomes binding and the principal has to resort to so-called "ironing" to construct an incentive-compatible mechanism (see, e.g., Myerson, 1981; Toikka, 2011). In the optimal mechanism from Section 2.3, bunching occurs even though assumption (2.16) explicitly rules out that the global incentive compatibility constraint is binding. Instead, the reason for bunching is the kink in the buyer's utility function.

Consequently, bunching is a feature of the optimal mechanism with explicit quantity and spending goals more generally. The utility function in (2.1) introduces kinks into the buyers utility by assuming a linear penalty when missing the goal. However, the essential feature is the presence of a kink in the buyers utility, not the specific functional form of the penalty.

Assume that the penalty for missing the quantity goal is given by $p_Q(x) > 0$ with $p'_Q(x) < 0$ (lower quantities are worse) for $x < \hat{x}$ and the penalty for missing the spending goal is $p_B(t) > 0$ with $p'_B(t) > 0$ for $t > \hat{t}$. Assume both penalty functions are weakly convex, which implies that the marginal penalty increases with the size of the deviation from the goal. For example, this includes the case of measuring the penalty for deviating from the quantity target in terms of $v(x)$. With weakly convex penalty functions, the result of Remark 2.1 continues to hold: The buyer will implement a mechanism whose outcomes only depend on the type of the most efficient supplier. The buyer then maximises

$$\int_{\underline{\theta}}^{\bar{\theta}} \left\{ \Big(v(x(\theta)) - t^*\big(x(\theta)\big) \Big) - \mathbf{1}_Q(\cdot)\, p_Q(x(\theta)) - \mathbf{1}_B(\cdot)\, p_B\Big(t^*\big(x(\theta)\big)\Big) \right\} f_1(\theta) d\,\theta,$$

where the incentive compatible transfer schedule $t^*(\cdot)$ is unchanged and given by (2.4). Appendix 2.B shows that, for constant $\mathbf{1}_Q$ and $\mathbf{1}_B$, the first-order condition can be written as

$$\begin{aligned} v'(x(\theta)) = \theta + \frac{F(\theta)}{f(\theta)} + \mathbf{1}_Q \underbrace{p'_Q(x(\theta))}_{<0} \\ + \mathbf{1}_B \underbrace{\left(\left(\theta + \frac{F(\theta)}{f(\theta)} \right) p'_B(t^*(\theta)) - \frac{\int_{\underline{\theta}}^{\theta} p''_B(t^*(z))\, t^{*\prime}(z)\, F(z)\, dz}{f(\theta)} \right)}_{>0}. \end{aligned} \tag{2.20}$$

Setting $\mathbf{1}_Q$ and $\mathbf{1}_B$ to one in (2.20) results in a discrete shift of the optimal quantity schedule: Missing the quantity goal leads to an upward, (globally) missing the spending goal to a downward shift of the optimal quantity. These discrete shifts of the optimal quantity schedule at the targets result in (partial) bunching regarding the quantity and total spending. As the generalised penalty functions do not change the optimal transfer schedule, it also remains true that a spending goal can impact the buyer's optimal behaviour even when they source from types for which they do not overspend. If the buyer suffers from overspending when contracting with efficient types, they curb their demand for inefficient types, even when their spending goal is not relevant for those types.

2.5 Conclusion

Auctions seem to be among the few forms of optimal procurement widely used in the real world. Buyers who aim to procure a fixed quantity often do so by auction. However, as seen in the case of the Austrian spectrum coverage Auction in 2020 (see Section 2.1), buyers sometimes also auction off fixed budgets to fix their spending. Moreover, most buyers do want to react to the marginal prices they face. This chapter has aimed to consider these observations not as isolated phenomena but as complementary perspectives on optimal procurement.

In the model of this chapter, the buyer benefits from direct consumption utility but suffers when the procured quantity falls short of the envisaged quantity or when spending exceeds the envisaged budget. If the buyer is unencumbered by penalties for missing their goals, the model corresponds to standard quasi-linear preferences. Thus, the buyer procures a higher quantity from more efficient suppliers. At the other extreme, the penalties for missing goals dominate the standard utility, and the buyer minimises total spending for buying a given quantity or maximises the amount procured with their fixed budget. These cases correspond to the preferences usually considered in the auction literature. In the optimal mechanism, the buyer engages in full bunching, i.e., they source the same quantity from all supplier types or spend the same amount on procurement from any type. Auctions with a fixed quantity or a fixed budget implement the optimal mechanism. If the buyer does have relevant goals but these goals do not entirely dominate the buyer's preferences, the buyer generally engages in partial bunching in terms of quantity and spending; there is a subset of adjacent (cost) types from which they source exactly their target quantity or on which they spend their target budget. Thus, fixed-quantity and fixed-budget auctions are not *the* but only *a* response to quantity and spending goals in procurement.

2.A Appendix: Proofs of Section 2.3

Section 2.3.3: Bounds on λ.— By the structure of the optimal mechanism from (2.14), the buyer switches to constant spending of $\hat{t}$ at θ_4. Thus, not accounting for the penalty on production costs, the marginal utility of additional procurement, i.e., the marginal utility in the first line of (2.10) is positive. Using the right-hand side of (2.15) for $v'(x)$, it follows that $\lambda > 0$. Consequently, $x^*_{BB}(\theta) < x^*_{BL;\mathbf{1}_B=0}(\theta)$. Second, the marginal utility of overspending is still negative when accounting for the penalty on production costs, i.e., the second line in (2.10) is positive. Hence, it must be that $\lambda < \bar{\lambda} = n\lambda_B\Big(F(\theta_4) - F(\theta_3) + \theta_4 f(\theta_4)\Big)$. It follows that $x^*_{BB}(\theta) > x^*_{BL;\mathbf{1}_B=1}(\theta)$.
Section 2.3.3: Optimal Interval Boundaries.— Partial differentiation of the relevant terms gives

$$\frac{\partial \Psi}{\partial x^*_{FB/BB}} = v'\left(x^*_{FB/BB}(\theta)\right) + \lambda_Q\,\mathbf{1}_Q - \left(\theta + \frac{F(\theta)}{f(\theta)}\right), \tag{2.A.1}$$

$$\frac{\partial x^*_{FB}}{\partial C} = \frac{1}{(1-F(\theta))^{n-1}}, \tag{2.A.2}$$

$$\frac{\partial x^*_{BB}}{\partial \lambda} = \frac{1}{n\,f(\theta)\;v''\left(x_{BB}(\theta)\right)}, \tag{2.A.3}$$

and

$$\frac{\partial x^*_{BB}}{\partial \theta_3} = \frac{\lambda_B\, f(\theta_3)}{f(\theta)\;v''\left(x_{BB}(\theta)\right)}. \tag{2.A.4}$$

Implicitly differentiating the condition $x^*_{FB}(\theta_3; C) = x^*_{BL;\mathbf{1}_B=1}(\theta_3)$ yields

$$\frac{\partial C}{\partial \theta_3} = -\frac{x^{*\prime}_{FB}(\theta_3) - x^{*\prime}_{BL}(\theta_3)}{\left.\frac{\partial x^*_{FB}}{\partial C}\right|_{\theta=\theta_3}}. \tag{2.A.5}$$

Implicit differentiation of the equation $x^*_{FB}(\theta_4; C) = x^*_{BB}(\theta_4; \lambda, \theta_3)$ results in

$$\frac{\partial \theta_4}{\partial \theta_3} = -\frac{\left.\frac{\partial x^*_{FB}}{\partial C}\right|_{\theta=\theta_4}\frac{\partial C}{\partial \theta_3} - \left.\frac{\partial x^*_{BB}}{\partial \theta_3}\right|_{\theta=\theta_4}}{x^{*\prime}_{FB}(\theta_4) - x^{*\prime}_{BB}(\theta_4)} \tag{2.A.6}$$

and

$$\frac{\partial \theta_4}{\partial \lambda} = \frac{\left.\frac{\partial x^*_{BB}}{\partial \lambda}\right|_{\theta=\theta_4}}{x^{*\prime}_{FB}(\theta_4) - x^{*\prime}_{BB}(\theta_4)}. \tag{2.A.7}$$

Last, implicit differentiation of the condition pinning down λ gives

$$\frac{d\lambda}{d\theta_3} = -\frac{\theta_4 \frac{\partial x^*_{BB}}{\partial \theta_3}\Big|_{\theta=\theta_4} + \frac{\int_{\theta_4}^{\bar{\theta}} \frac{\partial x^*_{BB}}{\partial \theta_3}(1-F(\theta))^{n-1}\, d\theta}{(1-F(\theta_4))^{n-1}} + t^{*\prime}_{BB}(\theta_4)\frac{\partial \theta_4}{\partial \theta_3}}{\theta_4 \frac{\partial x^*_{BB}}{\partial \lambda}\Big|_{\theta=\theta_4} + \frac{\int_{\theta_4}^{\bar{\theta}} \frac{\partial x^*_{BB}}{\partial \lambda}(1-F(\theta))^{n-1}\, d\theta}{(1-F(\theta_4))^{n-1}} + t^{*\prime}_{BB}(\theta_4)\frac{\partial \theta_4}{\partial \lambda}}, \tag{2.A.8}$$

where

$$t^{*\prime}_{BB}(\theta_4) = \theta_4 x^{*\prime}_{BB}(\theta_4) + (n-1)\,\frac{\int_{\theta_4}^{\bar{\theta}} x^*_{BB}(\theta)\,(1-F(\theta))^{n-1}\,d\theta}{(1-F(\theta_4))^n}\, f(\theta_4)$$

Section 2.3.4: Coinciding Bunching Intervals.— Consider the case where in the schedule from (2.18) there is θ_5 such that $x_{FB}(\theta_5) = \hat{x}$ with $\theta_5 \in (\theta_3, \theta_4)$. Such a case can occur given that the quantity schedules $x^*_{BB}(\theta)$ and $x^*_{FB}(\theta)$ do not directly depend on $\hat{x}$. For $\theta \to \theta_5^+$, the buyer engages in bunching regarding spending because $v'(x_{FB}(\theta)) - (\theta + F(\theta)/f(\theta)) + \lambda_Q - (\lambda/n + \lambda_B\, F(\theta_3))/f(\theta) > 0$ (i.e., the marginal utility of additional procurement under the regime $x^*_{BB}(\theta)$ is positive), and simultaneously $v'(x_{FB}(\theta)) + \lambda_Q - (1+\lambda_B)\,(\theta + F(\theta)/f(\theta)) < 0$ (i.e., the marginal utility of additional procurement under the regime $x^*_{BL}(\theta)$ is negative). At θ_5 both marginal utilities jump downward by λ_Q, as the buyer no longer benefits from lowering the penalty for missing their quantity goal when procuring more. Thus, the buyer will not procure more than $x_{FB}(\theta)$ for $\theta \to \theta_5^-$.

Case a: If the quantity goal is not too important, it still holds that $v'(x_{FB}) - (\theta + F(\theta)/f(\theta)) - (\lambda/n + \lambda_B\, F(\theta_3))/f(\theta) \geq 0$; the buyer (weakly) benefits from additional procurement as long as $t(\theta) < \hat{t}$. Thus, the buyer keeps spending constant while procuring more than $\hat{x}$. Therefore, the optimal quantity schedule is

$$x^*_{QB'}(\theta) = \begin{cases} x^*_{BL;\mathbf{1}_B=1,\mathbf{1}_Q=0}(\theta) & \text{for } \theta \leq \theta_3 \\ x^*_{FB}(\theta) \text{ with } x^*_{FB}(\theta_3) = x^*_{BL;\mathbf{1}_B=1,\mathbf{1}_Q=0}(\theta_3) & \text{for } \theta \in (\theta_3, \theta_4) \\ x^*_{BB;\mathbf{1}_Q=1}(\theta) & \text{for } \theta \geq \theta_4. \end{cases} \tag{2.A.9}$$

Case b: By contrast, for a more important quantity goal, $v'(x_{FB}) - (\theta + F(\theta)/f(\theta)) - \lambda/(n\, f(\theta)) - \lambda_B\, F(\theta_3)/f(\theta) < 0$, such that the buyer optimally spends less than $\hat{t}$

even for $\theta < \theta_5$, which, again would result first in a spending bunching interval and then overspending for more efficient types. This implies that it was not optimal to follow $x^*_{FB}(\theta)$ in the interval $[\theta_5, \theta_4]$ to begin with: If the buyer spends weakly less than $\hat{t}$ and overspends for more efficient types and there is a spending bunching interval for intermediate types, the optimal form of the quantity schedule is given by $x^*_{BB}(\theta)$ not $x^*_{FB}(\theta)$. Thus, this case collapses into the situation of (2.19).

Section 2.3.4: Spending Drops below $\hat{t}$ in the Quantity Bunching Interval.— The quantity schedule in (2.19) is only optimal if spending remains above $\hat{t}$ in the interval $[\theta_1, \theta_2]$. Instead, assume that there is $\theta_6 \in (\theta_1, \theta_2)$ such that $t^*(\theta_6)$ under the mechanism from (2.19). For $\theta \to \theta_6^-$, we cannot have $x > x_{FB}$. To prove this by contradiction, suppose that $x > x_{FB} > \hat{x}$. In that case, the optimal quantity schedule would be $x^*_{BL;\mathbf{1}_Q=0,\mathbf{1}_B=1}(\theta)$. However, $x^*_{BL;\mathbf{1}_Q=0,\mathbf{1}_B=1}(\theta) < \hat{x}$ by the construction of (2.11). Therefore, the buyer spends weakly less than $\hat{t}$. As the buyer overspends when sourcing from very efficient types by assumption (2.13), the optimal quantity schedule takes the shape of $x^*_{BB}(\theta)$. Given that $x^*_{BL;\mathbf{1}_Q=1,\mathbf{1}_B=1}(\theta) > x_{FB}$, $x^*_{BB;\mathbf{1}_Q=1}(\theta) > x_{FB}$. Thus, the optimal quantity schedule is determined by the position of $x^*_{BB;\mathbf{1}_Q=0}(\theta)$, i.e., the importance of the quantity goal.

Case a: If the quantity goal is not too important, $x^*_{BB;\mathbf{1}_Q=0}(\theta) > \hat{x}$, which, by assumption (2.12), also implies that $x^*_{BB;\mathbf{1}_Q=0}(\theta) > x_{FB}(\theta)$. As seen in Section 2.3.3, a situation where $x^*_{BB;\mathbf{1}_Q=0}(\theta) > x_{FB}(\theta) > x^*_{BB;\mathbf{1}_Q=0}(\theta)$ results in bunching regarding spending. Thus, the optimal quantity schedule is

$$x^*_{QB''}(\theta) = \begin{cases} x^*_{BL;\mathbf{1}_B=1,\mathbf{1}_Q=0}(\theta) & \text{for } \theta \leq \theta_1 \\ x^*_{FB2}(\theta) \text{ with } x^*_{FB2}(\theta_1) = x^*_{BL;\mathbf{1}_B=1,\mathbf{1}_Q=0}(\theta_1) & \text{for } \theta \in (\theta_1, \theta_6) \\ \hat{x} & \text{for } \theta \in (\theta_6, \theta_2) \\ x^*_{AB;\mathbf{1}_Q=1}(\theta) & \text{for } \theta \in [\theta_2, \theta_3] \\ x^*_{FB}(\theta) \text{ with } x^*_{FB}(\theta_3) = x^*_{BL;\mathbf{1}_B=1,\mathbf{1}_Q=1}(\theta_3) & \text{for } \theta \in (\theta_3, \theta_4) \\ x^*_{BB;\mathbf{1}_Q=1}(\theta) & \text{for } \theta \geq \theta_4, \end{cases} \tag{2.A.10}$$

where $x^*_{AB;\mathbf{1}_Q=1}(\theta)$ accounts for the fact the buyer overspends for $\theta \in [\theta_2, \theta_3]$ but does not do so for all more efficient types, i.e., it solves

$$v'\Big(x^*_{AB;\mathbf{1}_Q=1}(\theta)\Big) = \left(\theta + \frac{F(\theta)}{f(\theta)}\right) - \lambda_Q + \frac{\lambda_B\, F(\theta_1)}{f(\theta)} + \lambda_B\, \theta.$$

Similarly, $x^*_{BB;\mathbf{1}_Q=1}(\theta)$ needs to be adapted to reflect that there is no overspending in the interval (θ_1, θ_6); it solves

$$v'\left(x^*_{BB;\mathbf{1}_Q=1}(\theta)\right) = \left(\theta + \frac{F(\theta)}{f(\theta)}\right) - \lambda_Q + \frac{\lambda}{n\, f(\theta)} + \frac{\lambda_B\Big(F(\theta_3) - F(\theta_2) + F(\theta_1)\Big)}{f(\theta)}.$$

Case b: If the quantity goal is sufficiently important, $x^*_{BB;\mathbf{1}_Q=0}(\theta) \leq \hat{x}$ such that the situation is just as in Section 2.3.2 and the buyer engages in bunching regarding the quantity; the buyer sources $\hat{x}$ units until $x^*_{BB;\mathbf{1}_Q=0}(\theta)$ indeed results in procurement of more than $\hat{x}$ units. For more efficient types, the mechanism has the same structure as in Section 2.3.3. Consequently, we have

$$x^*_{QB'''}(\theta) = \begin{cases} x^*_{BL;\mathbf{1}_B=1,\mathbf{1}_Q=0}(\theta) & \text{for } \theta \leq \theta_1 \\ x^*_{FB2}(\theta) & \text{for } \theta \in (\theta_1, \theta_8) \\ x^*_{BB2;\mathbf{1}_Q=0}(\theta) & \text{for } \theta \in [\theta_8, \theta_7] \\ \hat{x} & \text{for } \theta \in \Big(\theta_7, \theta_2\Big) \\ x^*_{AB;\mathbf{1}_Q=1}(\theta) & \text{for } \theta \in [\theta_2, \theta_3] \\ x^*_{FB}(\theta) & \text{for } \theta \in (\theta_3, \theta_4) \\ x^*_{BB;\mathbf{1}_Q=1}(\theta) & \text{for } \theta \geq \theta_4 \end{cases} \tag{2.A.11}$$

with $x^*_{FB2}(\theta_1) = x^*_{BL;\mathbf{1}_B=1,\mathbf{1}_Q=0}(\theta_1)$, $x^*_{FB}(\theta_3) = x^*_{BL;\mathbf{1}_B=1,\mathbf{1}_Q=1}(\theta_3)$, and

$$v'\left(x^*_{BB2;\mathbf{1}_Q=0}(\theta)\right) = \left(\theta + \frac{F(\theta)}{f(\theta)}\right) + \frac{\lambda_2/n + \lambda_B\, F(\theta_1)}{f(\theta)}.$$

2.B Appendix: Proofs of Section 2.4

First-Order Condition for Generalised Penalty Functions.— Making the buyer's utility non-linear in t makes the buyer's maximisation problem somewhat more complicated. To make the buyer's problem a well-defined problem of variational calculus / optimal control, define $\check{X}(\theta) = -\int_{\theta}^{\bar{\theta}} x(z)\left(1 - F(z)\right)^{n-1} dz$, such that $\check{X}'(\theta) = \check{x}(\theta) = x(\theta)\big(1 - F(\theta)\big)^{n-1}$ corresponds to the expected quantity a seller of type θ can expect to sell. With these definitions in place, we have $x(\theta) = \check{x}(\theta)\left(1 - F(\theta)\right)^{1-n}$ and $t^*(\theta) = \theta\left(1 - F(\theta)\right)^{1-n}\check{x}(\theta) - \check{X}(\theta)\big(1 - F(\theta)\big)^{1-n}$, such that the buyer solves

$$\begin{aligned}
\max_{\check{X}(\theta)} \quad & \underbrace{\int_{\underline{\theta}}^{\bar{\theta}} \left\{ v\left(\frac{\check{x}(\theta)}{(1-F(\theta))^{n-1}}\right) - \left(\frac{\theta\,\check{x}(\theta)}{(1-F(\theta))^{n-1}} - \frac{\check{X}(\theta)}{(1-F(\theta))^{n-1}}\right)\right\} f_1(\theta) d\theta}_{\text{Direct consumption utility}} \\
& \underbrace{- \int_{\underline{\theta}}^{\bar{\theta}} \mathbf{1}_Q(\cdot)\, p_Q\left(\frac{\check{x}(\theta)}{(1-F(\theta))^{n-1}}\right) f_1(\theta) d\theta}_{\text{Penalty for missing quantity goal}} \\
& \underbrace{- \int_{\underline{\theta}}^{\bar{\theta}} \mathbf{1}_B(\cdot)\, p_B\left(\frac{\theta\,\check{x}(\theta)}{(1-F(\theta))^{n-1}} - \frac{\check{X}(\theta)}{(1-F(\theta))^{n-1}}\right) f_1(\theta) d\theta}_{\text{Penalty for missing spending goal}} \\
\text{s.t.} \quad & \check{x}'(\theta) \leq 0
\end{aligned}$$

as long as $\mathbf{1}_{Q/B} = \text{const}$. Solving the corresponding Euler-Lagrange equation and rewriting the solution in terms of $x(\theta)$ and $t(\theta)$, we get that the optimal quantity schedule for constant $\mathbf{1}_{Q/B}$ satisfies

$$\begin{aligned} v'(x(\theta)) = \theta + \Big(2 - v''(x(\theta))x'(\theta)\Big)\frac{f(\theta)}{f'(\theta)} \\ + \mathbf{1}_Q \left(p_Q'(x(\theta)) + p_Q''(x(\theta))\, x'(\theta)\, \frac{f(\theta)}{f'(\theta)}\right) \\ + \mathbf{1}_B \left(\theta\, p_B'\big(t^*(\theta)\big) + \Big(2p_B'\big(t^*(\theta)\big) + \theta\, p_B''\big(t^*(\theta)\big)\, t^{*'}(\theta)\Big)\, \frac{f(\theta)}{f'(\theta)}\right). \end{aligned} \tag{2.B.1}$$

Perceiving the first-order condition in (2.B.1) as a linear first-order differential equation in terms of $v''(\cdot)$, we get the formulation in (2.20).

3
Optimal, Simple, and Robust Auctions

Unpublished Working Paper

Abstract

This chapter analyses the procurement decision of a buyer who has quasi-linear utility but suffers when procuring too little or spending too much. They procure from one of several suppliers, each with private information on their constant unit costs that are independent draws from a commonly known distribution. A first-price auction, where the traded quantity depends on the winning bid, implements the optimal mechanism. These optimal auctions are demanding for the buyer: computing the optimal demand schedule is complex, and they might not possess the relevant information on the distribution of suppliers' costs. To address these challenges, this chapter also considers simple (i.e., easy-to-implement) and robust auctions that perform well irrespective of the distribution of types. This chapter compares two simple auctions: quantity and budget auctions. In a quantity auction, the buyer always sources the same quantity. In a budget auction, they source the highest possible quantity given their budget. Both designs are optimal if the respective goal is sufficiently important. If both goals are equally important, a not too risk-averse buyer prefers the budget auction. In a robust auction, the buyer maximises their expected utility for the worst possible distribution of cost types. Any demand schedule that procures the first-best quantity from the highest-cost type does so.

3.1 Introduction

Generally, buyers are unaware of the costs of their potential suppliers. For the buyer, this lack of knowledge is unfortunate because it means that they are unable to source at production costs. As a rule, they must pay higher prices instead. Their natural goal is to keep these higher prices as low as possible; they want to design an optimal procurement mechanism. When economic theorists consider this problem, they appeal to the revelation principle:[1] We can implement the outcome of any mechanism through a direct and truthful mechanism. In a direct, truthful mechanism, the buyer incentivises every supplier to reveal their costs. As these mechanisms are easy to analyse, theoretical economists usually content themselves with finding an optimal, truthful, direct mechanism. For all their mathematical elegance, it seems exceedingly rare that real-world buyers actually use direct-revelation mechanisms. Instead, they have used auctions at least two millennia before any economist stepped into the world to discover the revelation principle (see Chapter 1.3). As it turns out, auctions, although generally not truthful, are also optimal under certain conditions (see Myerson, 1981).

The previous chapter has taken the typical view of modern mechanism design. The buyer knows the probability distribution of suppliers' costs but not their actual costs. Chapter 2 derived a direct revelation mechanism that is optimal for the buyer when they have quasi-linear preferences, but they also suffer when they miss their quantity or spending goal: They have to endure a penalty when they buy too little or spend too much. This chapter adopts a more practical perspective by examining the use of auctions in this context.

As a first step, Proposition 3.1 shows that the buyer can implement the optimal mechanism through a first-price auction à la Hansen (1988). Every supplier submits

[1] A first version of the revelation principle appears in Gibbard (1973). In its "modern" version, it is traced back to a series of papers by Myerson (1979, 1982, 1986).

a unit price. The supplier bidding the lowest price wins the auction. They supply a quantity that is determined by their bid through a pre-announced demand schedule at the price they bid. Casual empiricism suggests that, although real-world buyers commonly use auctions, they do not use *this type of auction.* Instead, they primarily opt for fixed-quantity auctions, which fix the traded quantity independently of the winning bid. One explanation might be that this reflects their preferences: As Chapter 2 has shown, it is optimal to fix the traded quantity if the buyer's quantity goal is sufficiently demanding and important. If that is true, there is nothing more to investigate.

However, it might also be that real-world buyers do not implement optimal mechanisms. To design an optimal mechanism, the buyer must know the distribution of suppliers' potential costs and, based on that distribution, compute the optimal demand schedule, which can be a fairly complex task. The economic literature has recently taken up such considerations under the names of simple and robust mechanism design (for reviews, see Li, 2024 and Carroll, 2019, respectively). In that spirit, this chapter complements the perspective of optimal auctions with perspectives of simple and robust auctions.

There is no agreed-upon definition of what constitutes a "simple" mechanism. Intuitively, fixed-quantity auctions are definitely simple. The buyer only needs to announce the quantity they want to buy; they do not need to use any information on the distribution of supplier types. Moreover, this mechanism is even optimal when the quantity goal is sufficiently demanding and important. Consequently, it is plausible that a buyer relying on simple auctions would choose a fixed-quantity auction. There is another demand schedule in the class of potentially optimal auctions that shares the same properties: The buyer could also announce their target spending and buy as many units as possible at the winning price. This approach also does not require knowledge of the distribution of supplier types. As before, such a fixed-budget auction is optimal if the buyer has a very demanding and important spending target. Hence, if the primary concern is simplicity, the buyer could also conduct a fixed-budget auction. Of course, the quantity auction will do better to the degree that the quantity goal is important and vice versa. However, Proposition 3.2 shows that even a buyer who attaches a somewhat higher importance to their quantity goal (or who has purely quasi-linear preferences) prefers a budget auction as long as they do not suffer too much from variation in the procured quantity. In a fixed-budget auction, the buyer buys more when unit costs are low and less when unit prices are high. Thus, they can procure a higher expected quantity with the same expected spending. As the procured quantity varies in the fixed-budget auction, the buyer only benefits from this higher expected quantity if they can tolerate this variation.

Robust mechanism design is (among others) concerned with situations where the buyer cannot attach a probability to the winning supplier being of a given cost type; the buyer lacks the information to implement an optimal auction. Following a standard approach in the literature on robust mechanism design, this chapter assumes that the buyer evaluates mechanisms based on their expected utility under the worst-case distribution of types (maxmin expected utility) in such a scenario. Any mechanism that implements the first-best quantity for the least-efficient type is optimal under this robustness criterion (Proposition 3.3).

This chapter proceeds as follows. Section 3.2 revisits the model. Section 3.3 shows that a first-price auction with a variable quantity schedule can implement the optimal mechanism. Section 3.4 considers simple auctions and compares the relative performance of two easy-to-implement auctions: fixed-quantity and fixed-budget auctions. Section 3.5 derives optimal robust auctions that do not require information on the distribution of supplier types. Section 3.6 concludes.

3.2 Model

Buyer.— The buyer sources a quantity x from one supplier and pays that supplier a transfer of t. Their utility from this trade is

$$u(x,t) = \underbrace{\Big(v(x) - t\Big)}_{\text{Direct utility}} \underbrace{- \lambda_Q \, \mathbf{1}_{x \le \hat{x}} \, (\hat{x} - x) - \lambda_B \, \mathbf{1}_{t \ge \hat{t}} \, (t - \hat{t})}_{\text{Penalty functions: penalty for missing goals}}, \tag{3.1}$$

where $\hat{x}$ and $\hat{t}$ are their target quantity and their target spending, respectively. The buyer has a standard quasi-linear direct utility but suffers from buying too little or spending too much. The direct consumption utility $v(x)$ is an increasing and concave function with $\lim_{x \to 0^+} v'(x) = \infty$ and $\lim_{x \to \infty} v'(x) = 0$. The indicator functions take the value one if the procured quantity (total spending) is weakly below (weakly exceeds) its target $\hat{x}$ ($\hat{t}$) and are zero otherwise. Consequently, the buyer suffers to the degree λ_Q when their procurement falls short of the desired quantity and to the degree λ_B when they spend more than their envisaged budget. As $\lim_{x \to 0^+} v'(x) = \infty$, the buyer sources a strictly positive quantity.

Suppliers.— As in the traditional SIPV model (e.g., Krishna, 2010), there are n potential suppliers. Each supplier i has constant unit costs θ_i, which are private information of supplier i. It is common knowledge that the suppliers' unit costs are independent draws from the distribution $F(\theta)$ on $[\underline{\theta}, \bar{\theta}]$ with $0 < \underline{\theta} < \bar{\theta}$. Suppliers maximise their expected profits. A supplier of type θ with a revenue of t has profits of $t - \theta x$ when they produce x units. Their outside option has a value of zero.

Optimal Mechanism.— Assuming that the buyer sources from the supplier who claims to be of the most efficient type and only pays a transfer to that type, Chapter 2 has shown that a truthful, direct revelation mechanism solves

$$\max_{x(\theta)} \quad \int_{\underline{\theta}}^{\bar{\theta}} u(x(\theta), t^*(x(\theta))) \; f_1(\theta) \; d\theta \tag{3.2a}$$

$$\text{s.t.} \quad \frac{d}{d\theta} \, x(\theta) \Big(1 - F(\theta)\Big)^{n-1} \le 0, \tag{3.2b}$$

where $f_1(\cdot)$ denotes the density of the distribution of the most efficient supplier's type, and the optimal transfer is

$$t^*(\theta) = \theta\, x^*(\theta) + \int_\theta^{\bar{\theta}} x^*(z)\Big(1 - F(z)\Big)^{n-1} dz\, \Big(1 - F(\theta)\Big)^{1-n}. \tag{3.2c}$$

The optimal transfer from (3.2c) compensates suppliers at least for their production costs $\theta x^*(\theta)$ such that all supplier types participate in the tender. Additionally, it is the lowest transfer, which ensures that suppliers do not report being less efficient (local incentive compatibility). The mechanism is truthful if, additionally, suppliers do not report being more efficient than they actually are, which the global incentive compatibility constraint in (3.2b) guarantees.

3.3 Optimal Auctions

Analogously to the result of Dasgupta and Spulber (1989) for purely quasi-linear preferences (i.e., the special case where $\lambda_Q = \lambda_B = 0$) with a more general cost structure, Proposition 3.1 shows that the buyer can implement the optimal mechanism in a first-price auction.

The First-Price Auction Implementation.— In this first-price auction implementation, the buyer announces a quantity schedule $\tilde{x}(b)$, which specifies the number of units they buy if the per-unit price is b. Every potential supplier i simultaneously bids a per-unit price b_i. The supplier i^* who bids the lowest per-unit price supplies $x(b_{i^*})$ units at the unit price b_{i^*} they specified. This implementation corresponds to the sealed-bid auctions for the right to sell to a market with downward-sloping demand considered by Hansen (1988).

Implementing the Optimal Mechanism through a First-Price Auction.— Proposition 3.1 shows that the buyer can implement any solution to (3.2), for which the global incentive compatibility constraint (3.2b) is not binding, in such a first-price auction.

Proposition 3.1 (Dasgupta and Spulber, 1989: Implementation of Optimal Mechanism through First-Price Auction) *The buyer can implement any solution $x^*(\theta)$ to the optimisation problem in* (3.2)*, for which the global incentive compatibility constraint* (3.2b) *is not binding, as a symmetric pure-strategy Bayes-Nash equilibrium of the first-price auction implementation game with the quantity schedule*

$$\tilde{x}^*(b) = x^*(b^{*-1}(b)) \tag{3.3a}$$

with

$$b^*(\theta) = \frac{t^*(\theta)}{x^*(\theta)} = \theta + \frac{\int_\theta^{\bar{\theta}} x^*(z)\Big(1 - F(z)\Big)^{n-1} dz}{x^*(\theta)\Big(1 - F(\theta)\Big)^{n-1}}, \tag{3.3b}$$

where $b^*(\theta)$ is the equilibrium strategy of suppliers.

Proof. First, establish that $b^{*'}(\theta) > 0$, such that $b^*(\theta)$ is invertible, and we can define the quantity schedule in (3.3a). We have

$$b^{*'}(\theta) = -\frac{\int_\theta^{\bar{\theta}} x^*(z)\Big(1 - F(z)\Big)^{n-1} dz}{x^*(\theta)\Big(1 - F(\theta)\Big)^{n-1}} \frac{\frac{d}{d\theta} x^*(\theta)\Big(1 - F(\theta)\Big)^{n-1}}{x^*(\theta)\Big(1 - F(\theta)\Big)^{n-1}},$$

which is strictly positive if and only if $\frac{d}{d\theta} x^*(\theta)(1 - F(\theta))^{n-1} < 0$ and zero if $\frac{d}{d\theta} x^*(\theta)(1 - F(\theta))^{n-1} = 0$. Thus, the quantity schedule in (3.3a) and, by extension, the implementation game, is well-defined if and only if (3.2b) does not bind in optimum.

Second, if the suppliers bid according to (3.3b), the auction game indeed implements the optimal mechanism: As $b^{*'}(\theta) > 0$, the most efficient supplier wins the contract and supplies $\tilde{x}^*(b^*(\theta)) = x^*(b^{*-1}(b^*(\theta))) = x^*(\theta)$ units in exchange for a transfer of $x^*(\theta) \cdot b^*(\theta) = x^*(\theta) \cdot (t^*(\theta)/x^*(\theta)) = t^*(\theta)$.

Third, establish that suppliers bidding $b^*(\theta)$ is a symmetric equilibrium. If all other suppliers follow the strategy $b^*(\theta)$, a supplier of type θ_i chooses from the same set of possible outcomes as in the direct revelation game. If they bid $b^*(\theta_i)$, they reach the same expected profit as with truthful reporting in the direct revelation game; if they submit some other price b_i, they realise the expected profit that they would obtain by reporting the type $b^{*-1}(b_i)$ in the direct revelation game. As their optimal strategy in the direct revelation game is to report their true type θ_i by (3.2b) and (3.2c), their optimal strategy is to bid $b^*(\theta_i)$ in the auction implementation. ■

The buyer's utility function only enters the proof of Proposition 3.1 indirectly: It determines whether the global incentive compatibility constraint (3.2b) is binding in the optimum, but is otherwise irrelevant. As the difference between this model and the one in Dasgupta and Spulber (1989) is a change in the utility function, namely adding penalty terms to the buyer's preferences, Proposition 3.1 is a trivial extension of the result in Dasgupta and Spulber (1989) for quasi-linear preferences.

The Special Cases of Fixed-Quantity and Fixed-Budget Procurement.— For two extreme cases regarding the buyer's preferences, Chapter 2 has already established the optimal auction. When the buyer's quantity goal $\hat{x}$ is sufficiently demanding and they attach extreme weight to not missing the quantity goal ($\lambda_Q \to \infty$), they should engage in fixed-quantity procurement; their optimal demand schedule is

$$\tilde{x}_{QA}(b) = \hat{x}, \tag{3.4}$$

irrespective of the distribution of supplier types. Similarly, they should engage in fixed-budget procurement if their spending goal $\hat{t}$ is sufficiently tight and extremely important ($\lambda_B \to \infty$). They can implement this result by demanding

$$\tilde{x}_{BA}(b) = \hat{t}/b \tag{3.5}$$

units from a supplier winning the first-price auction with a bid of b. The optimality of these two quantity schedules is quite remarkable: The buyer does not need to know the distribution of the suppliers' types. They would not even benefit from having this knowledge. They do not even need to know the type space, i.e., the possible realisations of the suppliers' cost types, to implement the optimal mechanism. This is a much stronger result than the "traditional" results on the optimality of these auctions due to Myerson (1981) and Liu (2007). Under the preferences Myerson (1981) and Liu (2007) consider, auctions are only optimal when the auctioneer implements the optimal reserve price; to do so, they must know the distribution of types.

3.4 Simple Auctions

The Perspective of Simple Mechanism Design.— Undoubtedly, there is an inherent mental complexity in computing the optimal mechanism from Proposition 2.1 of Chapter 2 and, potentially, transforming the resulting quantity schedules into an auction according to (3.3) even when all relevant information is readily available. This is especially true if the buyer has both a quantity and a spending goal, in which case the optimal mechanism might require the separate optimisation of several parts of a piecewise function. In the face of such complexity, some management scholars argue that managers *should* resort to heuristics (e.g., Bettis, 2017) or accept satisfactory outcomes (e.g., Simon, 1955). In line with this argument, Camboni *et al.* (2025) show experimentally that theoretically superior but more complex procurement auctions might do worse than simpler alternatives. Even if one does not endorse the normative claim that managers should use heuristics, it seems reasonable to suggest that, empirically, they do.[2] Thus, it would be useful if mechanism designers identified the best mechanism for a particular purpose within the class of simple mechanisms. However, there is no commonly accepted formal

[2]In an extreme example, around 50% of U.S. CEOs report that their gut feeling is (very) important for their decisions on capital allocation Graham *et al.* (2015). Directly related to procurement, Turcic *et al.* (2024) show that the procurement decisions of the automobile producer BMW conform with theoretical considerations *within the confines of the contractual structure BMW uses.* However, the optimal structure of contracts is more complex.

definition of "simplicity" that allows delineating the class of simple mechanisms (Börgers, 2015, p. 165; Carroll, 2019, p. 159). To the degree that formal criteria exist, they have focused on the simplicity of mechanisms for the participants rather than the designer of the mechanism (see Li, 2024, for a review).

Fixed-Quantity and Fixed-Budget Procurement as Management Heuristics.— Hence, the goal of this section is not to find the best simple mechanism but to compare two simple mechanisms: fixed-quantity and fixed-budget first-price auctions, i.e., implementing the quantity schedules from (3.4) and (3.5) in a first-price auction of the type discussed in Section 3.3. Intuitively, they are simple because they only require the buyer to know a very basic part of their preferences: the procured quantity and the total spending that they target. Quantity and budget auctions seem reasonable as management heuristics or "good enough" mechanisms, as fixed-quantity and fixed-budget auctions are optimal if the buyer attaches extreme weight to their quantity and spending goal, respectively. Consequently, these auctions should not do too badly, even if the goals are somewhat less important to the buyer. This is also reflected in the fact that parts of the optimal quantity schedule from Chapter 2 resemble their outcomes more generally.

Which auction performs better depends on the relative weight the buyer attaches to not missing the quantity or spending goal. Proposition 3.2 shows that a buyer who does not suffer too much from variation in the procured quantity is better off under the budget auction when both goals are equally important. It follows that even a buyer who attaches slightly more weight to the quantity goal prefers the budget auction.

Equilibria in Fixed-Quantity and Fixed-Budget Auctions.— As a first step in the comparison, briefly consider suppliers' equilibrium behaviour in these auctions.

From routine steps, per-unit bids in the quantity and budget auction (with subscripts QA and BA, respectively) satisfy

$$b'_{QA}(\theta) = (n-1)\ (b_{QA}(\theta) - \theta)\ \frac{f(\theta)}{1-F(\theta)} \tag{3.6}$$

and

$$b'_{BA}(\theta) = (n-1)\ (b_{BA}(\theta) - \theta)\ \frac{f(\theta)}{1-F(\theta)}\ \frac{b_{BA}(\theta)}{\theta} \tag{3.7}$$

with $b_{QA/BA}(\bar{\theta}) = \bar{\theta}$ (see Krishna, 2010 for the quantity and Dastidar, 2008 for the budget auction), respectively.

Remark 3.1 (Lower Equilibrium Bids in Budget Auction) *Unit prices are lower in a first-price fixed-budget auction than in a first-price fixed-quantity auction whenever the winning bidder is of a type $\theta < \bar{\theta}$, i.e., $b_{BA}(\theta) < b_{QA}(\theta)$. When the winning bidder is of type $\bar{\theta}$, unit prices under the two auction designs are $\bar{\theta}$, i.e., $b_{BA}(\bar{\theta}) = b_{QA}(\bar{\theta}) = \bar{\theta}$.*

Proof. Applying Hôpital's rule to equations (3.6) and (3.7), we have $b'_{QA}(\bar{\theta}) = b'_{BA}(\bar{\theta}) = (n-1)/n < 1$. Thus, for (potentially, arbitrarily small) ε, $b_{QA}(\bar{\theta}-\varepsilon) > b_{BA}(\bar{\theta}-\varepsilon)$ as bids are steeper in the fixed-budget auction if bids are above costs. It follows that bids in the quantity auction are higher for all $\theta < \bar{\theta}$. Prove this by contradiction. If bids in the quantity auction were to fall below those in the budget auction for some $\theta < \bar{\theta}$, there would need to be a type $\check{\theta} < \bar{\theta}$ such that $b_{QA}(\check{\theta}) = b_{BA}(\check{\theta})$. As bids in the budget auction would be steeper at $\check{\theta}$ it would need to be the case that (in an environment of $\check{\theta}$) $b_{QA}(\theta) < b_{BA}(\theta)$ for $\theta > \check{\theta}$. ■

Intuitively, when lowering their bid, a bidder in the fixed-budget auction increases their probability of winning (as in the fixed-quantity auction) and, additionally, benefits from an increased traded quantity; the increased marginal benefit of lowering the bid results in lower unit prices in the fixed-budget auction.

Comparing Auction Designs: Defining Targets.— Comparing the relative performance of the two auction designs requires a plausible relation between the quantity and spending targets. When it comes to direct utility, a disproportionately low or high target spending would result in very low utility under a budget auction, as the buyer would procure far from the optimal quantity. Similarly, when it comes to the penalty for missing goals, choosing a very high spending target would favour the fixed-quantity auction, as it would not result in spending-related losses. Assume that the buyer has an implicit target price $\hat{p}$. They (implicitly) compute target spending from the target quantity (or vice versa) such that $\hat{x}\,\hat{p} = \hat{t}$. This assumption implies that they miss (at least) one goal whenever the unit price they pay is higher than $\hat{p}$, which ensures a consistent evaluation of outcomes between auction designs. Specifically, let

$$\hat{t} = \hat{x}\,\hat{p} \qquad \text{with} \qquad \hat{p} \in \Big[\mathbb{E}[b_{BA}(\theta_1)], \mathbb{E}[b_{QA}(\theta_1)]\Big] \tag{3.8}$$

where θ_1 denotes the winning bidder's cost type, i.e, the cost type of the most efficient supplier. Condition (3.8) assumes that the target price that is in the realm of prices they can expect under the two auction designs they consider.

Comparing Equilibrium Outcomes.— Liu (2007) and Liu and Parlour (2023) show that the key difference between quantity and budget auctions is that quantity-weighted prices are lower in a budget auction: in a budget auction, the buyer procures more when costs are low and vice versa. Consequently, they procure a higher (expected) quantity with the same (expected) budget as in the fixed-quantity auction. Remark 3.2 proves that the buyer procures a higher expected quantity under a budget auction when their target spending satisfies (3.8). This result drives the relative advantage of the budget auction.

Remark 3.2 (Higher Expected Quantity in Budget Auction) *If the spending target $\hat{t}$ satisfies* (3.8)*, a buyer auctioning off the budget $\hat{t}$ in a first-price budget auction in expectation procures more than $\hat{x}$ units, which is the quantity they procure in a comparable quantity auction.*

Proof. In expectation, the buyer procures $\mathbb{E}[t/b_{BA}(\theta_1)]$ units in the budget auction. By Jensen's inequality

$$\mathbb{E}[\hat{t}/b_{BA}(\theta_1)] \overset{(3.8)}{=} \hat{x}\,\hat{p}\,\mathbb{E}[1/b_{BA}(\theta_1)] \overset{(JI)}{\geq} \hat{x}\,\frac{\hat{p}}{\mathbb{E}[b_{BA}(\theta_1)]} \overset{(3.8)}{\geq} \hat{x},$$

where Jensen's inequality is strict for any non-degenerate distribution $F(\theta)$.[3] ■

Comparing Direct Utility.— Remark 3.2 directly implies that a risk-neutral buyer (who benefits from trade) has higher direct utility under the budget auction: By assumption (3.8), their budget in the budget auction is weakly lower than their expected spending in the quantity auction. Nonetheless, they procure a weakly higher expected quantity. However, the buyer with the quasi-linear preferences from (3.1) is not risk-neutral in that sense: They do not suffer from variation of their spending, but they do suffer when the procured quantity varies. This puts the budget auction at a disadvantage because it fixes spending at the cost of variation in the procured quantity. Intuitively, they are only better off under the budget auction if they do not suffer too much from this variation. Formally, consider the buyer as paying to participate in a lottery over quantities. The quantity auction gets them $\hat{x}$ units with certainty, whereas the distribution of the procured quantity in the budget auction is $G_{BA}(x) = 1 - F_1(b_{BA}^{-1}(\hat{t}/x))$, where $F_1(\cdot)$ is the distribution of the lowest of n draws from F. Denoting the certainty equivalent of a lottery G

[3] Consequently, the budget auction also procures a higher expected quantity if bids under the designs were identical (for example, under a second-price rule) as long as the distribution of supplier types is not degenerate.

when the utility from consumption is $v(x)$ by $c(G, v)$, the buyer has weakly higher quasi-linear utility in the budget auction, if and only if

$$c(G_{BA}, v) \geq v(\hat{x}) - \underbrace{(\hat{x}\,\mathbb{E}[b_{QA}(\theta_1)] - \hat{t})}_{\geq 0 \text{ by (3.8)}}. \tag{3.9}$$

To have higher direct utility in the budget auction, the buyer must not be too risk-averse in the sense that their certainty equivalent for the "budget auction gamble" is high enough to beat the certain pay-off of the "quantity auction gamble", accounting for the lower costs of participating in the budget auction. Requiring a high enough certainty equivalent for the relevant distribution is equivalent to requiring a sufficiently low Arrow-Pratt measure of absolute risk aversion for the relevant quantity levels (Mas-Collel *et al.*, 1995, p. 191).

Comparing Overall Utility.— The overall difference between the expected utility in the fixed-budget auction, $\mathbb{E}[U_{BA}]$, and the expected utility in the fixed-quantity auction, $\mathbb{E}[U_{QA}]$, is

$$\mathbb{E}[U_{BA}] - \mathbb{E}[U_{QA}] = \underbrace{\left(\mathbb{E}[U^D_{BA}] - \mathbb{E}[U^D_{QA}]\right)}_{\equiv \Delta U^D} + \underbrace{\left(\mathbb{E}[U^P_{BA}] - \mathbb{E}[U^P_{QA}]\right)}_{\equiv \Delta U^P},$$

where the superscripts D and P denote direct utility and the (negative) utility from penalties for missing the goals, respectively. Condition (3.9) has already established when the buyer prefers the budget auction in terms of direct utility, i.e., $\Delta U^D \geq 0$. The buyer prefers the budget auction if, additionally, they suffer from lower penalties for missing their goals under the budget auction. By construction, the buyer does not miss their quantity goal in the quantity auction and does not overspend in the budget auction. However, if they pay a unit price above $\hat{p}$, they incur a penalty for missing a goal in either case; if prices are too high, they overspend in the quantity auction, and they procure too little in the budget auction. Formally, the difference in the penalty-related utility, ΔU^P, is

$$\Delta U^P = -\int_{\hat{\theta}_{BA}}^{\bar{\theta}} \lambda_Q \left(\underbrace{\hat{x} - \hat{t}/b_{BA(\theta)}}_{>0} \right) f_1(\theta)\, d\theta + \int_{\hat{\theta}_{QA}}^{\bar{\theta}} \lambda_B \left(\underbrace{\hat{x}\, b_{QA}(\theta) - \hat{t}}_{>0} \right) f_1(\theta)\, d\theta,$$

where the buyer misses of of their goals sourcing from a less efficient type than $\hat{\theta}_{QA}$ and $\hat{\theta}_{BA}$, respectively, i.e., $b_{QA}(\hat{\theta}_{QA}) = b_{BA}(\hat{\theta}_{BA}) = \hat{p}$. Evidently, the quantity auction becomes more attractive in terms of penalties as the importance of the quantity goal, λ_Q, increases. Analogously, an increase in λ_B shifts the comparison in favour of the budget auction. If the buyer suffers sufficiently from missing their goals, the penalty-related utility component dominates their quasi-direct utility. Thus, regarding penalties, they prefer a quantity or budget auction if λ_Q or λ_B is high enough, respectively.[4]

Equally Important Goals.— Thus, an interesting comparison of the auction designs concerns a buyer who attaches equal weight to both goals. Intuitively, such a buyer suffers equally from procuring one unit too little and spending the (target) price of one unit too much. Formalising this intuition, a buyer considers both goals equally important if $\lambda_Q = \tilde{p}\,\lambda_B$, where $\tilde{p}$ is the price the buyer uses to evaluate the costs of one unit. There are two plausible candidates for $\tilde{p}$. First, the buyer could use the ex-ante expected price, i.e., $\tilde{p} = \mathbb{E}[b(\theta_1)]$. Second, they could use the price they actually pay in the given situation, as it reflects how the buyer could actually shift disutility between the two dimensions; in that case, we would have $\tilde{p} = b(\theta_1)$.[5] In both cases, the relevant unit prices could be those from the quantity or budget auction. Allowing both (and all prices in between), a buyer attaches equal weight to both goals if their preferences satisfy

$$\lambda_Q = \tilde{p}\,\lambda_B \qquad \text{with} \qquad \tilde{p} \in \Big[\mathbb{E}[b_{BA}(\theta_1)], \mathbb{E}[b_{QA}(\theta_1)]\Big] \tag{3.10}$$

[4]This is, assuming that $F(\theta)$ is not degenerate. If it is, they never miss a goal (see Proposition 3.2).

[5]In that case, λ_Q varies with the winning supplier's type and the mechanism for price determination. Consequently, it is inconsistent with the analysis in Chapter 2 and Proposition 3.1, which treats λ_Q as a constant.

or

$$\lambda_Q = \tilde{p}(\theta_1)\,\lambda_B \qquad \text{with} \qquad \tilde{p}(\theta_1) \in [b_{BA}(\theta_1), b_{QA}(\theta_1)]. \tag{3.11}$$

Advantage of the Budget Auction.— A buyer whose goals are equally important suffers from lower penalties in the budget auction. It follows that they prefer the budget auction for an arbitrary, non-degenerate distribution of suppliers' costs if they are not too risk-averse in the sense of (3.9) (Proposition 3.2).

Proposition 3.2 (Advantage of the Budget Auction) *Consider a buyer whose utility is given by* (3.1) *with targets satisfying* (3.8). *If the buyer is not too risk-averse in the sense of* (3.9) *and attaches equal weight to both goals in the sense of* (3.10) *or* (3.11), *they are strictly better off when they conduct a fixed-budget auction of a budget* $\hat{t}$ *than when they buy* $\hat{x}$ *in a fixed-quantity auction for any non-degenerate distribution of supplier types. For a degenerate distribution, the buyer is indifferent.*

Proof. The previous text shows that, under the conditions of Proposition 3.2, the buyer has higher quasi-linear utility in the budget auction, i.e., $\Delta U^D > 0$ for a non-degenerate distribution. Appendix 3.A establishes that these conditions also imply that they suffer from lower penalties in the budget auction, i.e., $\Delta U^P > 0$. A degenerate distribution has a mass point at $\bar{\theta}$ (the highest possible type). Hence, $b_{QA}(\bar{\theta}) = b_{BA}(\bar{\theta}) = \bar{\theta} = \hat{p}$. Consequently, both auction designs are equivalent, as they procure $\hat{x}$ units at total costs of $\hat{t}$ with certainty. ■

The reason for lower penalties in the budget auction is identical to that for higher quasi-linear utility: As Remark 3.2 shows, the budget auction offers a better relation between the expected procured quantity and expected spending.

3.5 Robust Auctions

The Informational Requirements of Implementing the Optimal Mechanism.— Implementing the optimal auction is not only mentally complex but also informationally demanding. As a rule, the buyer must know the precise distribution $F(\theta)$ of the suppliers' production costs to define the demand schedule from (3.3a). Because this arguably is a tall order in real applications (see, e.g., the evidence in Kim, 2025), "there is an ethos in much of the mechanism design community that *realistic* mechanisms should not be finely tuned to parametric assumptions, such as probability distributions of values [in the present case: unit costs]" (Carroll, 2019, p. 159, own emphasis). From a more practical perspective, the question then is how the buyer can do well even when they cannot assign probabilities to the winning supplier being of a specific type, i.e., when they operate under "Knightian uncertainty" (see Knight, 1921, pp. 19–20).[6]

The Perspective of Robust Mechanism Design.— From the viewpoint of economic theory, such a perspective makes the buyer's choice of a procurement mechanism a problem of what is called robust mechanism design in the particular flavour of what Carroll (2019) calls "robustness to distributions". In this literature, the designer of the mechanism evaluates the performance of mechanisms based on their performance under the worst-case distribution of types, as they cannot compute expected utility. In game-theoretic terms, they commit to a mechanism that maximises their pay-off given that an adversarial nature designs the distribution to minimise their pay-off after observing the designer's move. The most natural candidate for the buyer's pay-off is their expected utility, which gives rise to the

[6]Alternatively, the buyer might have *some* knowledge of the distribution of supplier types. For example, they might know that the suppliers' types follow one of several distributions (see, e.g., Bose *et al.*, 2006, Carroll, 2017, Che and Zhong, 2024 or, in a different context, Wolitzky, 2016).

so-called maxmin-utility approach (e.g., Auster, 2018; Bose *et al.*, 2006; Carrasco *et al.*, 2018, 2019; Che and Zhong, 2024).

Maxmin-Utility-Optimal Auctions.— Finding a robust first-price auction in the case at hand amounts to solving

$$\max_{\tilde{x}(b)} \min_{F(\theta)} \mathbb{E}[u(\tilde{x}, F)], \tag{3.12}$$

where $u(\tilde{x}, F)$ denotes the buyer's utility when they implement the demand function $\tilde{x}(b)$ and supplier types are drawn from the distribution $F(\theta)$. Proposition 3.3 describes a solution to (3.12).

Proposition 3.3 (Maxmin-Optimal Auctions) *Any weakly-decreasing quantity schedule $\tilde{x}(b)$ that (a) induces increasing bids in equilibrium and (b) only procures units that result in weakly positive marginal utility is maxmin-optimal if $\tilde{x}(\bar{\theta}) = x_{FB}(\bar{\theta})$, where $x_{FB}(\theta)$ is the first-best quantity procured from a supplier of type θ. This includes a quantity auction of $x_{FB}(\bar{\theta})$ units and a budget auction with a budget of $\bar{\theta}\, x_{FB}(\bar{\theta})$.*

Proof. To verify this equilibrium, consider what happens when the buyer announces a weakly downward-sloping demand schedule. The buyer suffers from higher bids given that $u'(b) = (v'(\tilde{x}(b)) - b)\tilde{x}'(b) - x(b) < 0$ if the buyer benefits from all units, i.e., $v'(\cdot) \geq b$. If less efficient types submit higher bids in equilibrium, the worst-case distribution puts all probability mass on the least efficient type $\bar{\theta}$. The buyer's best response is to implement the first-best quantity, $x_{FB}(\bar{\theta})$, for that type. Being in Bertrand competition, suppliers truthfully report their costs. Consequently, the buyer is indifferent between all weakly downward-sloping demand schedules that share the same quantity $x(\bar{\theta})$ if they induce increasing bids for a non-degenerate distribution. The description of equilibrium bids in equations (3.6) and (3.7) shows that the quantity and budget auctions satisfy this condition. ■

Alternative Notions of Robustness.— The maxmin criterion is not very discerning. For example, according to Proposition 3.3, both a quantity auction of $x_{FB}(\bar{\theta})$ units and a budget auction with a budget of $\bar{\theta}\, x_{FB}(\bar{\theta})$ are maxmin-optimal auctions. However, a buyer with a demanding quantity goal, who attaches high importance to reaching that goal, should clearly opt for a quantity auction; the quantity auction would implement the optimal mechanism for most plausible distributions of supplier types and not only for the worst-case distribution. Being overly pessimistic, the maxmin criterion does not even consider other distributions than the worst-case distribution. Consequently, some authors sharpen the maxmin criterion by requiring that the mechanism is undominated; there should be no other mechanism that performs "better in all plausible scenarios" (Börgers *et al.*, 2025, p. 3). Under this intuitive definition of dominance, quantity and budget auctions are the only undominated mechanisms for a buyer who attaches sufficient weight to their target quantity or spending, respectively.[7]

However, most of the literature responds to the problem that "the maxmin utility approach is too conservative to be useful" (Zhang, 2022, p. 3) by dismissing the maxmin-utility criterion altogether and turning to the minmax-regret criterion instead (e.g., Bergemann and Schlag, 2008, 2011; Koçyiit *et al.*, 2024; Zhang, 2022). This criterion replaces utility in (3.12) by the negative of the buyer's regret; thus, nature maximises regret, whereas the buyer minimises it. Regret is the difference between the utility the designer of the mechanism could have realised if they knew the winning supplier's type and the utility they realise under the mechanism they implement.[8] Consequently, the adversarial nature strikes a balance between two considerations. As with maxmin-utility, nature should put all probability weight on

[7]The values of λ_Q and λ_B that are sufficient for fixed-quantity and fixed-budget procurement being optimal depend on the class of relevant distributions.

[8]A few authors, such as Kasberger and Woodward (2025), also consider *loss* as target function, i.e., the difference between the buyer's expected utility had they known the *distribution* of types and the expected utility they realise.

the least efficient type to minimise the buyer's expected utility. However, nature also has an incentive to put probability weight on very efficient types, as this increases the utility the buyer could have realised.

The buyer does not influence the first-best utility they could have realised. Thus, they face the standard problem of maximising their expected utility for a given – in this case, adversarial – distribution of types (see Bergemann and Schlag, 2008). Consequently, if the quantity or spending goal is sufficiently important, the minmax-regret optimal auction is a quantity or budget auction because they are the expected utility maxmising mechanism no matter the distribution of types.

Outside these two extreme cases, finding a minmax-regret optimal auction is more challenging than in most settings considered in previous literature for two reasons. First, the traded quantity is endogenous rather than fixed (in contrast to, e.g., Bergemann and Schlag, 2008, 2011; Koçyıit *et al.*, 2024; Zhang, 2022). Thus, the buyer announces not only a price but a quantity-transfer tuple. Second, the buyer faces several suppliers (as opposed to Bergemann and Schlag, 2008, 2011) that behave strategically. This strategic behaviour cannot be reduced to truthful bidding as in the second-price settings in Zhang (2022) or Koçyıit *et al.* (2024) because the buyer prefers first-price auctions (see Remark 2.1). Hence, suppliers' equilibrium strategy depends on nature's distribution choice, complicating the construction of a minmax-optimal auction.

If bidding behaviour were independent of the distribution of types, we could construct a minmax-optimal auction as follows: The buyer chooses a mechanism that keeps their regret constant for all types. Nature, being indifferent between all distributions, chooses a distribution that makes this mechanism a best response. This approach fails in the case at hand because nature's distribution choice alters the buyer's regret in a given situation by changing the winning supplier's bid. Furthermore, the differential equation that describes equilibrium bids does not have a general closed form solution, making it difficult to account for the induced changes in bids. Considering a truthful direct revelation mechanism instead of a first-price auction resolves this problem. However, even then, a minmax-regret optimal mechanism is not readily available. Appendix 3.B discusses nature's distribution design choice and shows that there is no equilibrium, in which nature chooses a continuous distribution $F(\theta)$.

3.6 Conclusion

This chapter has considered a buyer who has quasi-linear preferences but also suffers from procuring too little or spending too much. The buyer sources from one of several suppliers, each with privately known constant unit costs drawn independently from a commonly known distribution.

Generally, the buyer can implement the optimal procurement mechanism in a first-price auction. Each supplier bids a unit price, the lowest bid, and the traded quantity is determined from the winning bid according to a pre-announced demand schedule. To design these auctions, the buyer must know the distribution of suppliers' unit costs and go through the potentially complex task of computing the optimal quantity schedule. The buyer might not be able or willing to make these computations and might not possess the relevant information. In light of this, this chapter has explored auction designs that are simple for the buyer and robust auctions, which do not require information on the distribution of supplier types.

Two of the potentially optimal quantity schedules seem particularly simple: In a quantity auction, the buyer always sources the same quantity. In a budget auction, they source the highest possible quantity with their given budget. Both auction designs are optimal if the respective goal is very demanding and sufficiently important. This chapter has shown that a buyer who attaches equal weight to both goals prefers a budget over a quantity auction.

If the buyer does not know the distribution of supplier types, the literature on robust mechanisms suggests that they should rank mechanisms according to the expected utility they realise under the worst-case distribution of types. In that case, they are indifferent between any quantity schedule that demands the first-best quantity from the least efficient type.

3.A Appendix: Proof of Proposition 3.2

Difference in Penalties for Non-Degenerate Distributions of Types.— By Remark 3.1, we have that $b_{QA}(\theta) > b_{BA}(\theta)$ for $\theta < \bar{\theta}$ and, therefore, $\hat{\theta}_{BA} > \hat{\theta}_{QA}$. Thus, the penalty-related utility in the fixed-quantity auction is bounded from above by

$$\begin{aligned}\mathbb{E}[U^P_{QA}] &\overset{\text{Rem. 3.1}}{<} -\int_{\hat{\theta}_{BA}}^{\bar{\theta}} \lambda_B \left(\hat{x}\, b_{QA}(\theta) - \hat{t}\right) f_1(\theta)\, d\theta \\ &\overset{(3.8)}{=} -\lambda_B\, \hat{x}\, (1 - F_1(\hat{\theta}_{BA})) \left(\mathbb{E}[b_{QA}(\theta_1)|\theta_1 > \hat{\theta}_{BA}] - \hat{p}\right),\end{aligned}$$

If $\tilde{p}$ is given by (3.10), we have

$$\begin{aligned}\mathbb{E}[U^P_{BA}] &\overset{\substack{\text{Rem.}\\ 3.1}}{>} -\int_{\hat{\theta}_{BA}}^{\bar{\theta}} \lambda_Q \left(\hat{x} - \frac{\hat{t}}{b_{QA}(\theta)}\right) f_1(\theta)\, d\theta \\ &\overset{\substack{\text{(JI)}\\ (3.8)}}{>} -\lambda_B \hat{x}(1 - F_1(\hat{\theta}_{BA})) \left(\frac{\tilde{p}}{\mathbb{E}[b_{QA}(\theta_1) \mid \theta_1 > \hat{\theta}_{BA}]}\right) \left(\mathbb{E}[b_{QA}(\theta_1) \mid \theta_1 > \hat{\theta}_{BA}] - \hat{p}\right).\end{aligned}$$

From these bounds, it follows directly that $\mathbb{E}[U^P_{BA}] > \mathbb{E}[U^P_{QA}]$ if $\tilde{p} < \mathbb{E}[b_{QA}(\theta_1)|\theta_1 > \hat{\theta}_{BA}]$, which holds given that $\hat{\theta}_{BA} > \underline{\theta}$ and $\tilde{p}$ satisfies (3.10).

If $\tilde{p}$ is given by (3.11), a lower bound for penalty-related utility in the budget auction is

$$\begin{aligned}\mathbb{E}[U^P_{BA}] &\overset{(3.11)}{\geq} -\int_{\hat{\theta}_{BA}}^{\bar{\theta}} \lambda_B\, b_{QA}(\theta) \left(\hat{x} - \frac{\hat{t}}{b_{BA}(\theta)}\right) f_1(\theta)\, d\theta \\ &\overset{(3.8)}{=} -\lambda_B\, \hat{x}\, (1 - F_1(\hat{\theta}_{BA})) \left(\mathbb{E}[b_{QA}(\theta_1)|\theta_1 > \hat{\theta}_{BA}] - \hat{p}\, \mathbb{E}\left[\frac{b_{QA}(\theta_1)}{b_{BA}(\theta_1)} \middle| \theta_1 > \hat{\theta}_{BA}\right]\right).\end{aligned}$$

Thus, a sufficient condition for $\mathbb{E}[U^P_{BA}] > \mathbb{E}[U^P_{QA}]$ is that $\mathbb{E}\left[\frac{b_{QA}(\theta_1)}{b_{BA}(\theta_1)} \middle| \theta_1 > \hat{\theta}_{BA}\right] \geq 1$, which holds by Remark 3.1.

3.B Appendix: Minmax-Regret Optimal Mechanism

Necessary Condition for Best Response of Buyer.— Assume the buyer has has quasi-linear preferences. The quantity schedule $x(\theta)$ they announce as a best response to nature's distribution choice equalises their marginal utility and the winning bidder's virtual costs (see equation (2.6) in Chapter 2), i.e.,

$$v'(x(\theta)) = \theta + \frac{F(\theta) - F(\underline{\theta})}{f(\theta)}.$$

Importantly, they procure the first-best quantity from the most efficient type which has a positive density, as $F(\theta) - F(\underline{\theta}) = 0$.

Nature's Distribution Choice.— If $F(\theta)$ is continuous, nature's distribution choice solves the optimal control problem

$$\max_{f(\theta)} \quad \int_{\underline{\theta}}^{\bar{\theta}} \Bigg\{ \underbrace{u^*(\theta)}_{\text{First-Best Utility}} - \underbrace{\left(v(x(\theta)) - \left(\theta + \frac{F(\theta) - F(\underline{\theta})}{f(\theta)} \right) x(\theta) \right)}_{\text{Actual Utility}} \Bigg\} f_1(\theta) d\theta$$

$$\text{s.t.} \qquad F'(\theta) = f(\theta)$$

$$1 = \int_{\underline{\theta}}^{\bar{\theta}} f(\theta)\, d\theta$$

$$f(\theta) \geq 0,$$

where $f_1(\theta) = n(1-F(\theta))^{n-1} f(\theta)$, $u^*(\theta) = v(x_{FB}(\theta)) - \theta x_{FB}(\theta)$, and $f(\theta)$ is nature's control variable. A continuous adversarial distribution is inconsistent with the buyer playing a best response to such a distribution. Prove this by contradiction.

Defining $r(\theta)$ as the term in the curly brackets and, nature's first-order condition requires that for all θ with a positive density

$$\begin{aligned} &\Big(u^*(\theta) - v\big(x(\theta)\big) + \theta\, x(\theta)\Big) \Big(1 - F(\theta)\Big)^{n-1} \\ &\quad + \int_{\underline{\theta}}^{\theta} r(z)\,(n-1)\Big(1 - F(z)\Big)^{n-2} f(z) - x(z)\Big(1 - F(z)\Big)^{n-1} dz = \mu, \end{aligned} \tag{3.B.2}$$

where μ is a constant. In equilibrium, nature cannot benefit from removing probability mass somewhere and moving it elsewhere. The value of μ is such that nature overall places a probability mass of one. A continuous $F(\theta)$ requires that (3.B.2) holds for an interval of adjacent types. Consider the lowest type, $\tilde{\theta}_1$, on which nature places a positive probability mass, i.e., for which (3.B.2) must hold. Playing a best response to the distribution nature chooses, the buyer procures the first-best quantity from that type. This implies that nature does not place any probability mass on the type $\tilde{\theta}_1 + \varepsilon$ with $\varepsilon \to 0^+$: Because the buyer procures the first-best quantity, we have $u^*(\tilde{\theta}_1) - v\big(x(\tilde{\theta}_1)\big) + \tilde{\theta}_1\, x(\tilde{\theta}_1) = \frac{d}{d\theta}\Big(u^*(\theta) - v\big(x(\theta)\big) + \theta\, x(\theta)\Big)\Big|_{\theta=\tilde{\theta}_1} = 0$. Furthermore, by the same token, $r(\tilde{\theta}_1) = 0$. Thus, the left-hand side of (3.B.2) decreases by $x_{FB}(\tilde{\theta}_1)(1 - F(\tilde{\theta}_1))^{n-1}$ for a marginal increase in θ. Consequently, the left-hand side of (3.B.2) is smaller than μ for $\tilde{\theta}_1 + \varepsilon$ and nature does not place any density on that type (they would benefit from moving probability mass to $\tilde{\theta}_1$). Consider the next most efficient type $\tilde{\theta}_2 > \tilde{\theta}_1$, to which nature attaches positive density. As there is only a singleton with positive density in the interval $[\underline{\theta}, \tilde{\theta}_2)$, $F(\tilde{\theta}_2) = F(\tilde{\theta}_1) = F(\underline{\theta})$ by continuity. Thus, the buyer sources the first-best quantity from type $\tilde{\theta}_2$ and, again, nature does not attach any weight to type $\tilde{\theta}_2 + \varepsilon$. The same argument continues to apply. As the buyer also cannot implement a degenerate distribution, given that this would result in zero regret, there is no equilibrium, in which nature plays a continuous distribution.

4
Auctioning off Budgets in Multi-Winner Procurement Auctions

A version of this chapter was originally published in The RAND Journal of Economics (Vol. 56(4), pp. 655–667) under the title "Auctioning Off Budgets in Procurement" (doi: 10.1111/1756-2171.70028). Minor revisions have been made for consistency with this dissertation. The use of the published article and these modifications are permitted under the terms of the CC BY 4.0 licence. The published version is also included in Appendix A.

Abstract

This chapter investigates a multi-unit pay-as-bid procurement auction where the auctioneer fixes total spending and maximises the quantity procured with their predetermined, secret budget. Previous literature has analysed fixed-quantity auctions, where the traded quantity is fixed but unknown to the bidders. Compared to such auctions, budget auctions lower the auctioneer's costs by introducing an additional interaction between a bidder's bids; bidders not only weigh a higher profit margin on a unit against a lower probability of supplying that unit; a higher margin on some unit also reduces the probability that the budget suffices to procure more units from the bidder.

4.1 Introduction

In 2020, the Austrian telecommunications regulatory authorities (TKK and RTR) used a reverse auction to procure mobile network coverage for remote areas that had previously lacked an adequate network infrastructure (RTR, 2020). In the pay-as-bid auction, each mobile network operator submitted a set of bids, each of which specified the total subsidy they demanded for providing coverage to a certain number of municipalities (TKK, 2019). The auctioneer selected the combination of the different bidders' bids that maximised the number of municipalities to which coverage would be provided, subject to the constraint that the total subsidy did not exceed a secret, predetermined budget (TKK, 2019).

The Austrian spectrum auction can be considered the real-world equivalent of the model in this chapter. The model describes a novel strategic effect in bidding behaviour. Due to the auctioneer's budget cap, there is an interaction between a bidder's bids on different units: A high bid on some unit reduces the residual budget. Thus, fewer of the bidder's bids can be considered. Such a link between bids on different units does not have a correspondence in previous literature.

This chapter derives the symmetric equilibrium in a budget auction when bidders have the same cost function and share a stochastic belief about the unknown budget of the auctioneer. Previous literature (e.g., Holmberg, 2009; Pycia and Woodward, 2026) has considered the case where the auctioneer fixes the traded quantity, which is unknown to the bidders. In a pay-as-bid procurement auction in such a setting, when bidders raise their bids, they trade off the additional margin on a particular unit (price effect) with a decreased probability of realising the margin on this unit (quantity effect). By contrast, increasing the profit margin on some unit also hurts the chances of realising the profit margin on *other* units in the budget auction. Higher bids decrease the residual budget and, therefore, the probability that bids on higher quantities can be served without exceeding the auctioneer's budget (budget effect). Consequently, the auctioneer's costs in a budget auction are strictly lower than in a comparable fixed-quantity auction, i.e., a fixed-quantity auction with the same distribution of realised demand in terms of quantities. The cost advantage of the budget auction is most pronounced when traditional fixed-quantity auctions suffer from soft competition, i.e., when there are few bidders whose costs exhibit strong diseconomies of scale.

This chapter proceeds as follows. Section 4.2 discusses related literature. Section 4.3.1 first explains the basic model setup for the fixed-quantity and the budget auction. For ease of exposition, it then briefly discusses the symmetric equilibrium in the fixed-quantity auction translating the model from Pycia and Woodward (2026) to the procurement context. Afterward, I characterise the symmetric equilibrium in the strategically more complex budget auction. Section 4.4 shows that the budget auction comes at a lower cost to the auctioneer than a comparable fixed-quantity auction. Section 4.5 discusses some of this chapter's limitations before Section 4.6 concludes.

4.2 Related Literature

Traditionally, auction theory assumes that the traded quantity is fixed in advance. In the "single-unit", or rather, "single-winner" context, the analysis of auctions with an endogenous traded quantity dates back at least to Hansen (1988). Dasgupta and Spulber (1989) demonstrate that variable-quantity auctions are optimal under plausible assumptions, both in cases with a single winner and multiple winners. However, in contrast to typical real-world auctions or budget auctions, implementing these optimal auctions requires the auctioneer to possess extensive knowledge regarding the bidders' cost structure. Budget auctions, as a specific form of auctions with a variable traded quantity, have been discussed in the single-winner context by Dastidar (2008), Deck and Wilson (2008), and Liu and Parlour (2023).

However, this chapter is better understood against the backdrop of the literature explicitly characterising equilibria in multi-unit (multi-winner) auctions. Burkett and Woodward (2020b) are a notable outlier in this literature because they consider a *discrete* uniform-price auction with *uncertainty over individual bidder types*. The rest of this literature, including this chapter, rather follows Klemperer and Meyer (1989) in analysing multi-winner auctions with a continuous traded quantity and bidders with an identical, known cost structure. Specifically, Klemperer and Meyer (1989) consider a uniform-price procurement auction in which the bidders share a belief about the random, downward-sloping demand function. They explicitly characterise the potentially infinite number of symmetric equilibria. They show that the equilibrium bidding functions lie in a band that bounds them from above and below. If one were to adapt their model such that demand is completely inelastic, i.e., the traded quantity is fixed in advance, this would correspond to the uniform-price equivalent of the model of a pay-as-bid fixed-quantity auction in this chapter. In such a model, the infinitely many symmetric equilibria can be arbitrarily unattractive to the auctioneer. These equilibria are called the "low-revenue equilibria" of the uniform price auction. Their occurrence was first observed by Wilson (1979).

For the pay-as-bid auction, Holmberg (2009) derives the symmetric equilibrium for the reverse auction with completely inelastic demand. Just like Klemperer and Meyer (1989), he assumes that bidders have an identical, known cost function and share a belief over the distribution of demand. However, for his model to work, he requires that demand surpasses the suppliers' production capacity with a positive probability, which limits the applicability of his results. Pycia and Woodward

(2026) relax this assumption in the context of a forward auction. Their model forms the fixed-quantity baseline to which I compare the budget auction in this chapter.

Pycia and Woodward (2026) also consider an endogenous traded quantity in their analysis of random reserve prices. Random reserve prices punish high bids by increasing the probability that the increased bid will not be considered. This reinforces the quantity effect and is, therefore, possible in both fixed-quantity and budget auctions. Thus, the insights into secret reserve prices complement the analysis in this chapter, which focuses on a novel strategic effect deterring high bids by reducing the probability that bids on other units are accepted.

For the uniform-price auction, some more papers consider a variable traded quantity besides Klemperer and Meyer (1989). These models address the issue of low-revenue equilibria by allowing the auctioneer to restrict the traded quantity after observing the bids (e.g., Back and Zender, 2001; Lengwiler, 1999; McAdams, 2007).[1] Like in this chapter, the threat to reduce the traded quantity if bids are too unattractive to the auctioneer prompts the bidders to reduce bid shading. However, the auction designs from this literature can only be applied in the real world under exceptional circumstances. Consider McAdams (2007), for example. In his model, the auctioneer adjusts supply *at will* ex-post, i.e., does not make *any* commitment regarding the traded quantity, and the bidders reveal their type *truthfully* as a consequence.

The auctioneer could exploit this design by learning the bidders' types without trading any significant quantity and then extracting rents from the bidders in trade outside the auction. According to Rothkopf *et al.* (1990), the bidders' unwillingness to reveal their type is one of the main reasons for the scarcity of second-price auctions in the real world. This concern is mitigated in the pay-as-bid setting

[1] Other possible solutions have been explored, for example, by Kremer and Nyborg (2004), Damianov (2005), and Burkett and Woodward (2020a).

considered here. Furthermore, the auctioneer in this chapter does commit to a budget, although the specific amount is not disclosed. Bergemann and Hörner (2018) and Allen *et al.* (2024) suggest that such opacity regarding parts of the auction design is not unusual in real-world auctions.

This raises the question of how the auctioneer can credibly commit when the bidders cannot observe this commitment. Tillio *et al.* (2016) argue that credible commitment is possible if the auctioneer is indifferent between the available alternatives. However, this seems unlikely in most applications and is certainly not the case for the budget auction. This suggests other mechanisms exist by which the auctioneer can make their private commitment credible. These include trust, as in the case of the Austrian coverage auction, and trusted intermediaries, such as eBay, for secret reserve prices on the platform. Theoretically, there are also trustless solutions. For example, the auctioneer could publish the budget in an encrypted document before and the encryption key after the auction. However, it seems unlikely that procuring entities would opt for such a solution. Equally theoretically, uncertainty over the budget could also stem from tying the budget to a random but ex-post-verifiable outcome, for example, proceeds from future auctions of emission allowances. Why the auctioneer might *want* to implement an uncertain budget is discussed in Section 4.5.

4.3 Model

4.3.1 Setup

The auctioneer wants to procure multiple units of a homogeneous good using a pay-as-bid auction. The main claim of this chapter is that the auctioneer is better off when doing so using a budget auction than when using a fixed-quantity design, which has been analysed by Holmberg (2009) and Pycia and Woodward (2026).

Bidders.— For both auction designs, assume that all of the $n \geq 2$ bidders have the same marginal cost $c(x) > 0$ for producing the x-th unit with $0 < c'(x) < \infty$.[2] The bidders know each others' costs, whereas the auctioneer does not. The auction rules require each bidder to specify a weakly positive, twice differentiable, finite, and increasing bidding function $b_i(x)$, which denotes the marginal payment bidder i demands for providing their x-th unit. I assume that the submitted bid functions are smooth in the sense that they have *finite* first and second derivatives for $x > 0$.[3] The inverse of $b_i(x)$ is denoted by $\beta_i(\cdot)$. The auctioneer pays bidder i $B_i(x) = \int_0^{x_i} b_i(y)\, dy$ when buying x_i units from them. I refer to $B_i(x)$ as bidder i's cumulative bids. The pay-off of bidder i when assigned quantity x is

$$\pi_i(x; B_i) = B_i(x) - \int_0^x c(y)\, dy.$$

Fixed-Quantity Auction.— In the fixed-quantity auction, the auctioneer procures X units, whatever the costs of doing so are. The auctioneer minimises the total payments to bidders subject to the constraint that X units are procured. In line with previous literature, assume that the bidders do not know the quantity the auctioneer wants to procure but believe that it is distributed on some interval $[0, \bar{X}]$ according to the distribution function $G(X)$ with $G'(X) = g(X) > 0$ and $\bar{X} < \infty$. This setup for the fixed-quantity auction is the reverse auction equivalent of the model from Pycia and Woodward (2026).

Budget Auction.— In the budget auction, the auctioneer intends to procure as many units of the good as possible with the fixed budget A. Given the submitted bids, the auctioneer solves the following optimisation problem to determine the

[2] It only makes sense for the auctioneer to conduct a multi-unit auction, i.e., source from multiple bidders, if marginal costs are increasing on some relevant interval.

[3] This is important because in most cases, the symmetric equilibrium in the budget auction can only be obtained numerically, which is impossible with an infinite slope.

quantity x_i they buy from bidder i:

$$\begin{aligned} \max_{x_1,\dots,x_n} \quad & \sum_{i=1}^{n} x_i \\ \text{s.t.} \quad & \sum_{i=1}^{n} B_i(x_i) \le A \\ & \forall i:\ x_i \ge 0. \end{aligned}$$

Assume that each bidder submits bids at least up to the quantity that the auctioneer wants to assign to them. The first-order conditions of this problem imply that in an interior solution with $x_i > 0$ for all bidders, all bidding functions must have the same value at their respective chosen quantity. Refer to this unique value as the market (clearing) price. Consequently, with increasing bids, the auctioneer procures all units for which the demanded price is lower or equal to the market price.

The auctioneer's budget is unknown to the bidders. They only share a common belief that the auctioneer's budget follows a twice-differentiable cumulative distribution $F(A)$ on $[0, \bar{A}]$ with density $f(A) > 0$, $\bar{A} < \infty$ and $f(\bar{A}) < \infty$, as well as $f'(\bar{A}) < \infty$. The bidders' beliefs are rational in the sense that $A \in [0, \bar{A}]$.

4.3.2 Symmetric Equilibrium in the Fixed-Quantity Auction

Pycia and Woodward (2026) have extensively analysed the fixed-quantity auction. Proposition 4.1 summarises their result on equilibrium behaviour.

Proposition 4.1 (Pycia and Woodward, 2026: Equilibrium in the Fixed-Quantity Auction) *The symmetric equilibrium in a pay-as-bid fixed-quantity auction, in which n suppliers with costs $c(x)$ believe that the demanded quantity is distributed according to the distribution function $G(\cdot)$ on $[0, \bar{X}]$, is described by*

$$b(x) = c(x) + [1 - G(nx)]^{\frac{1-n}{n}} \int_{x}^{\bar{X}/n} c'(y)\, [1 - G(ny)]^{\frac{n-1}{n}}\, dy.$$

Bids are strictly above cost for all but the last unit and equal cost for the last unit.

To better understand the differences between the fixed-quantity and the budget auction regarding equilibrium behaviour, the following lays out the crucial steps in the derivation of Proposition 4.1.

As bids are increasing, there is a unique market clearing price p, which ensures that demand equals supply, i.e., $\sum_i \beta_i(p) = X$, and all units for which the bids are below p are supplied. The quantity x_i supplied by player i is implicitly defined by $p = b_i(x_i)$. Following Pycia and Woodward (2026), the probability that player i supplies weakly less than x units given some set of strategies $\boldsymbol{b} = (b_i, \boldsymbol{b_{-i}})$ is

$$\text{Prob.}(x_i \leq x; \boldsymbol{b}) = G\Big(x + \sum_{j \neq i} \beta_j(b_i(x))\Big).$$

Bidder i winning some quantity x implies a market clearing price of $b_i(x)$, which means that the other bidders supply $\sum_{j \neq i} \beta_j(b_i(x))$ units. Consequently, overall demand must not exceed $x + \sum_{j \neq i} \beta_j(b_i(x))$ if player i is to win x units or less. As shown by Pycia and Woodward (2026), this probability distribution allows us to write the expected profits of bidder i as

$$\max_{b_i(x)} \quad \mathbb{E}[\pi_i] = \int_0^{\bar{x}_i} [b_i(x) - c(x)] \left[1 - G\Big(x + \sum_{j \neq i} \beta_j(b_i(x))\Big)\right] dx, \tag{4.1}$$

where $\bar{x}_i = \inf\{x : G(\cdot) = 1\}$. The intuitive interpretation of equation (4.1) is that the profit margin on a particular unit must be weighted by the probability with which it is realised. Due to the pay-as-bid nature of the auction, the profit margin on some quantity x is realised whenever the bidder wins x or more units.

As there is (abstracting away from the monotonicity constraint on bids) no interaction between the bids on different quantities, the first-order condition to the bidder's problem (4.1) is obtained by point-wise maximisation of the integrand. Holmberg (2009) and Pycia and Woodward (2026) derive the first-order condition as

$$\underbrace{\Big(1 - G(X)\Big)}_{\text{Price effect}} - \underbrace{\Big[b_i(x) - c(x)\Big]\, g(X) \sum_{j \neq i} \beta_j'(b_i(x))}_{\text{Quantity effect}} = 0, \tag{4.2}$$

where the market clearing condition $x_i + \sum_{j \neq i} \beta_j(b_i(x)) = X$ has been used.[4]

The first-order condition in equation (4.2) demonstrates the trade-off between price and quantity effect. The first term represents the positive effect of raising bids. In all cases where the bidder wins more than x units, they benefit from additional revenue if they increase $b_i(x)$ (price effect). The second term shows that an increase in the bid on unit x reduces the probability that the corresponding profit margin is realised (quantity effect). This effect constitutes the competition aspect of the auction. If a bidder demands a higher payment for some unit, the auctioneer considers more of the other bidders' bids first. This reduces the probability that the bidder gets to supply the unit in question.

Given that the quantity assigned to bidder i is strictly increasing in total demand X, equation (4.2) allows us to gain further insights into optimal bidding behaviour. As long as total demand is lower than $\bar{X}$, the first term in equation (4.2) is positive, such that the first-order condition can only be satisfied if the corresponding bids are above marginal costs. By the same argument, for the highest quantity ever won, which is the quantity $\bar{x}_i$ won when $X = \bar{X}$, it must be that the corresponding bid equals marginal cost, as $1 - G(\bar{X}) = 0$. Intuitively, the price effect of increasing the bid on the last unit is zero, as the additional margin is only realised when $X = \bar{X}$. By contrast, the risk of losing the (profitable) unit when increasing the bid remains.

Thus, we obtain the terminal condition $b_i(\bar{x}_i) = c(\bar{x}_i)$. Taken together with symmetry, i.e., $b_i(x) = b_j(x)\,\forall i, j$, this allows us to write the symmetric equilibrium, based on equation (4.2), as

$$b(x) = (n-1)\Big[1 - G(nx)\Big]^{\frac{1-n}{n}} \int_x^{\bar{x}} [1 - G(ny)]^{-1/n}\, g(ny)\, c(y)\, dy,$$

[4]Due to both technical reasons and the fact that Holmberg (2009) considers electricity markets, he analyses a somewhat different case, where demand exceeds production capacities with a positive probability. The model of Pycia and Woodward (2026) needs to be adapted to the case of a procurement auction.

where $\bar{x}_i = \bar{x} = \bar{X}/n$. Integration by parts results in the formulation of Proposition 4.1.

4.3.3 Symmetric Equilibrium in the Budget Auction

Next, consider the symmetric equilibrium in the budget auction.

Bidder's Problem.— Given a strategy profile $\boldsymbol{B} = (B_i, \boldsymbol{B_{-i}})$, for bidder i to win x units or less, it must be that the budget suffices at most to pay for the x units of bidder i and the $\sum_{j\neq i} \beta_j(b_i(x))$ units that the other bidders will supply in this situation, i.e.,

$$\text{Prob.}(x_i \leq x; \boldsymbol{B}) = F\Big(B_i(x) + \sum_{j\neq i} B_j\big(\beta_j(b_i(x))\big)\Big).$$

As before, the bidder's expected profit is given by weighting the profit margin on some unit x with the probability that the bidder wins at least x units. The problem of bidder i is

$$\max_{B_i(x)} \quad \mathbb{E}[\pi_i] = \int_0^{\bar{x}_i} \underbrace{[b_i(x) - c(x)]\left[1 - F\Big(B_i(x) + \sum_{j\neq i} B_j\big(\beta_j(b_i(x))\big)\Big)\right]}_{\equiv\Psi(B_i, b_i, x)} dx \tag{4.3a}$$

$$B(0) = 0 \tag{4.3b}$$

$$B_i(\bar{x}_i) = \underbrace{\bar{A} - \sum_{j\neq i} B_j\big(\beta_j(b_i(\bar{x}_i))\big)}_{\equiv\phi(b_i, \bar{x}_i)}. \tag{4.3c}$$

The bidder's best response problem comes with two boundary conditions. First, by construction, cumulative bids start at zero. Second, the interval of relevant quantities is endogenously determined by the considered bidder's bids. As bids are increasing, the highest quantity a bidder can win is the quantity the bidder wins when the auctioneer auctions off the highest possible budget $\bar{A}$.

Strategic Effects.— A closer look at equation (4.3a) reveals that the bid on some unit x is chosen according to three considerations:

1. Its effect on the profit margin on unit x (price effect).

2. Its effect on the probability of realising the profit margin on unit x (quantity effect).

3. Its effect on the probability of winning units $y > x$ (budget effect).[5]

The price effect is reflected in the first bracket of the integrand in (4.3a). The quantity effect is described by the second summand in the argument of the distribution function. It reflects competition in the sense that it determines the order in which the different bidders' bids are considered by the auctioneer. As was seen in Section 4.3.2, the trade-off between price and quantity effect characterises the equilibrium in the fixed-quantity auction. The budget effect, which is reflected in the $B_i(x)$ in the argument of the distribution function, has no correspondence in the fixed-quantity auction. To make the distinction between the effects clear, think of bids and cumulative bids as independent variables. In this case, demanding a higher total compensation does not contribute to profits. However, it reduces the residual budget and thus makes it less likely that the budget suffices to buy further units from the bidder. Interestingly, in contrast to the quantity effect, this negative effect of raising bids is not a competition effect. It is not about the relative attractiveness of the different bidders' bids. Rather, higher cumulative bids of a bidder shift the probability of winning higher quantities downward independently of what the other bidders do.

Equilibrium Bidding Behaviour.— As the bidder's problem from (4.3a) includes an interaction between the bids on different units, the first-order condition for this variational calculus problem is (see, e.g., Chiang, 1992)

$$\frac{\partial \Psi}{\partial B_i} - \frac{d}{dx}\frac{\partial \Psi}{\partial b_i} \stackrel{!}{=} 0. \tag{4.4}$$

[5]Technically, this formulation is somewhat imprecise, as a *single* bid does not influence a bidder's cumulative bids and, therefore, does not influence the probability of winning $y > x$ units.

Roughly speaking, $\partial\Psi/\partial B_i$ represents the budget effect, $\partial\Psi/\partial b_i$ is the balance of price and quantity effect, and the d/dx accounts for the relation between bids and cumulative bids. The relevant derivatives can be found in Appendix 4.A.1. Using symmetry, i.e., $b(x) = b_i(x) = b_j(x)\ \forall i, j$, the Euler equation emerges as the third-order differential equation

$$b''(x) = \frac{[b'(x)]^2}{b(x)[b(x) - c(x)]} \Bigg[3b(x) - \frac{(n-2)}{(n-1)}\, c(x) - \frac{c'(x)}{b'(x)}\, b(x) + n\, \frac{f'(nB(x))}{f(nB(x))}\, [b(x) - c(x)]\, \frac{[b(x)]^2}{b'(x)} \Bigg].^6 \tag{4.5}$$

The two boundary conditions in (4.3b) and (4.3c) eliminate two degrees of freedom. The initial condition (4.3b) removes the first degree of freedom. The second piece of information lies in the terminal condition from (4.3c). For the symmetric case, this implies

$$B(\bar{x}) = \frac{\bar{A}}{n}. \tag{4.6}$$

However, equation (4.6) only pins down the level of cumulative bids at the last unit but does not tell us what the last unit is. Appendix 4.A.2 shows that the corresponding transversality condition implicitly defines the last unit through $b_i(\bar{x}_i) = c(\bar{x}_i)$. Equivalently, for the symmetric case, where $\bar{x} \equiv \bar{x}_i = \bar{x}_j$, we have

$$b(\bar{x}) = c(\bar{x}). \tag{4.7}$$

Equation (4.7) says that the last unit should be chosen in such a manner that all opportunities to realise additional profits are exhausted; i.e., there is no positive margin on the last unit. Put differently, of the three strategic effects discussed earlier, only the price effect favours higher bids. However, as the weight bidders attach to the profit margin on the last unit is zero, it is absent for the last unit.

[6] Normally, the Euler equation is a second-order differential equation. However, simultaneously determining all bidders' mutual best responses adds a degree of freedom.

The only effect active on the last unit is the quantity effect. Thus, the last unit should be offered at marginal cost.

Finally, Appendix 4.A.2 shows that the smoothness assumption on $b(x)$ implies

$$b'(\bar{x}) = \frac{n-1}{2n-1}\, c'(\bar{x}). \tag{4.8}$$

Any increasing continuation of the bidding function beyond $\bar{x}$ stabilises the equilibrium.[7] Proposition 4.2 summarises the findings on the symmetric equilibrium.

Proposition 4.2 (Equilibrium in the Budget Auction) *The symmetric equilibrium in a pay-as-bid budget auction, in which n suppliers with costs $c(x)$ believe that the budget is distributed according to the distribution function $F(\cdot)$ on $[0, \bar{A}]$, is described by equations* (4.3b) *and* (4.5)–(4.8)*. Bids are increasing and strictly above cost for all but the last unit. Bids equal costs for the last unit.*

In addition to the previously discussed, Proposition 4.2 claims that bids are above marginal cost and satisfy the monotonicity constraint imposed in the model setup. The corresponding proofs and the second-order conditions are relegated to appendices 4.A.3 and 4.A.4, respectively. Generally, there is no closed form for the symmetric equilibrium.

Example.— The following example, which allows for a closed-form representation, will later serve to illustrate the cost advantage of the budget auction. Consider the case where $n = 2$, $c(x) = \alpha x + \gamma$, and $A \sim U[0, \bar{A}]$ with $\alpha, \gamma, \bar{A} > 0$. The symmetric equilibrium is

$$b(x) = \frac{\sqrt{36\gamma^2 + 60\alpha\bar{A}} + 9\gamma}{15} + \frac{\alpha}{3}\, x.$$

The second-order condition for this example is discussed in Appendix 4.A.4.

[7] Off-equilibrium quantities are only relevant, if a deviating player reduces the maximum quantity they win. This would either require $b_i(\bar{x}_i) > c(\bar{x}_i)$, or $b_i(\bar{x}_i) = c(\bar{x}_i) < b_j(\bar{x}_j) = c(\bar{x}_j)$, where $i \neq j$ and $\bar{x}_j > \bar{x}_i$. In either case, the bidder could profitably win further units.

4.4 Cost Advantage of the Budget Auction

Auctioning off a secret budget instead of an unknown quantity introduces an additional consideration into the bidders' decision. Under both auction designs, when raising bids, the bidders need to weigh additional profits on a unit (price effect) against a lower probability of realising them (quantity effect). However, in the budget auction, increasing the bid on some quantity has the additional adverse effect of making it less likely to win higher quantities, as the auctioneer runs out of budget "earlier" (budget effect). Therefore, intuitively, the budget auction should be more attractive to the auctioneer.

Testing this hypothesis requires fully specifying the distributions in which bidders believe under the two auction designs, respectively. Luckily, the bidders' knowledge of equilibrium bids allows them to seamlessly move between expectations over total spending and the procured quantity. Given their equilibrium bids in the budget auction, the bidders can map every potential budget into a quantity, which the auctioneer procures with that budget in total. Thus, the bidders can directly translate the probability distribution over budgets, which describes their beliefs in a budget auction, into a "corresponding" probability distribution over the number of units procured by the auctioneer. This corresponding distribution describes the bidders' belief in a comparable fixed-quantity auction.

For the case of symmetric bidders, the number of units X the auctioneer procures in a budget auction with a given budget is implicitly given as $A = n\, B^B({}^{X}\!/_{n})$ (where the superscript B denotes the budget auction). Thus, the probability that the auctioneer procures less than X units is equivalent to the probability that the budget does not exceed $nB^B({}^{X}\!/_{n})$. Therefore, the bidders believe that the auctioneer's demand follows the distribution

$$\tilde{G}(X) \equiv F\Big(n\, B^B({}^{X}\!/_{n})\Big) \tag{4.9}$$

in the corresponding fixed-quantity auction. These corresponding auctions can also be generalised to bidders who do not bid the same in equilibrium because they are asymmetric. In this case, we can write the budget needed to procure a given total quantity through a budget auction as a function of the quantity x_i supplied by bidder i. We then get

$$\tilde{G}(X) \equiv F\Big(B_i^B(x_i) + \sum_{j \neq i} B_j^B\Big(\beta_j^B\Big(b_i^B(x_i)\Big)\Big)\Big) \tag{4.10}$$

with the implicit definitions $A = B_i^B(x_i) + \sum_{j \neq i} B_j^B(\beta_j^B(b_i^B(x_i)))$ and $x_i = X - \sum_{j \neq i} \beta_j^B(b_i^B(x_i))$. This gives $\tilde{g}(X) = f(\cdot)\frac{\partial A}{\partial x_i}\frac{\partial x_i}{\partial X} = f(\cdot)b_i^B(x_i)$ because a marginal change in the procured quantity must be scaled by the corresponding payment. Comparing equilibrium bids in a budget auction with the corresponding fixed-quantity auction yields two results. First, directly formalising the intuition from the previous section, when fixing the other bidders' bids, a bidder's best response is to bid higher in in the fixed-quantity auction than in the budget auction (Proposition 4.3). Second, for symmetric bidders, this also entails lower equilibrium bids in the budget auction (Proposition 4.4).

Proposition 4.3 (Lower Best Response Bids in Budget Auction) *Consider a fixed-quantity auction with demand normalised according to equation* (4.10). *Fix the bids of bidders $j \neq i$ in the fixed-quantity auction at the level of their equilibrium bids in the corresponding budget auction, i.e., $b_j(x) = b_j^B(x)$. Then, some bidder i's best response is to bid $b_i(x) > b_i^B(x)$ for $x < \bar{x}_i$.*

Proof. Consider the best response $b_i^B(x)$ of bidder i in the budget auction. It has them bidding truthfully on the last unit they receive (see Appendix 4.A.2). With reference to equation (4.A.2), this means that $\frac{\partial \Psi}{\partial b_i^B}\Big|_{x_i = \bar{x}_i} = 0$. This formalises the intuition that the optimal bid equalises price and quantity effect because the budget effect is absent for the last unit. However, the bidder's best response has them bid below the level that would balance these two effects for all other units, i.e., $\forall x < \bar{x}_i : \frac{\partial \Psi}{\partial b_i^B} > 0$. Formally, the bidder's best response entails bidding above costs for these units (see Appendix 4.A.3). Thus, there is a budget effect in the sense that $\forall x < \bar{x}_i : \frac{\partial \Psi}{\partial B_i^B} < 0$, as can be seen from equation (4.A.1). Therefore, by the first-order condition (4.4), $\forall x < x_i : \frac{d}{dx}\frac{\partial \Psi}{\partial b_i^B} < 0$. It follows that along the equilibrium path $\forall x < \bar{x}_i : \frac{\partial \Psi}{\partial b_i^B} > 0$. The normalisation of demand implies that $\frac{\partial \Psi}{\partial b_i^B}$ is equivalent to the derivative in the first-order condition of the bidder's problem in the fixed-quantity auction if evaluated at $b_i(x) = b_i^B(x)\,\forall i$. Intuitively, the normalisation ensures that the density attached to every unit is the same under both auction designs such that price and quantity effects are identical. If the bidder's problem in the fixed-quantity auction is well behaved, i.e., concave in b_i, the level of their best response bids is above $b_i^B(x)$. ■

Although Proposition 4.3 demonstrates the incentive wedge between comparable budget and fixed-quantity auctions, it does not have direct implications for the equilibrium outcomes. Proposition 4.4 entails those at the cost of restricting attention to symmetric equilibria. It shows that equilibrium bids in a budget auction, denoted by $b^B(x)$, are lower than the equilibrium bids, $b^{FQ}(x)$, in a comparable fixed-quantity auction, where demand is normalised according to equation (4.9).

Proposition 4.4 (Lower Equilibrium Bids in the Budget Auction) *If demand is normalised according to* (4.9)*, bids in the symmetric equilibrium are lower in the budget auction than in the comparable fixed-quantity auction:*

$$b^B(x) \leq b^{FQ}(x) \quad \forall x \in [0; \bar{x}],$$

where the inequality is strict for $x \in [0, \bar{x})$ and the superscripts FQ and B denote the fixed-quantity and budget auction, respectively.

The result of Proposition 4.4 is very strong in that it postulates that *every single unit* is less expensive in the budget auction.

The equality of bids for the last unit follows from the fact that the maximum number of units an individual bidder supplies is identical between the corresponding auctions by construction. As bidders do not have a positive margin on the last unit under both auction designs, the bids on the last unit are identical in the corresponding auctions.[8] The proof of the remainder of Proposition 4.4 is relegated to Appendix 4.B. A sufficient condition for higher equilibrium bids in the fixed-quantity auction is that $\frac{\partial \Psi}{\partial b_i}\Big|_{B_i=B_j=B^B, b_i=b_j=b^B} > 0$ for $x < \bar{x}$. This means that in a hypothetical situation without budget effect, bids should be higher than in the equilibrium of the budget auction. As the budget effect pulls toward lower bids, this is the case. Formally, the proof of this is analogous to the one in Proposition 4.3.

Example.— Let us revisit the example from Section 4.3.3. Equilibrium bids in the budget auction with $n = 2$, $c(x) = x$, and $A \sim U[0,1]$ are $B^B(x) = {}^1\!/\!{}_6\; x^2 + {}^{\sqrt{60}}\!/\!{}_{15}\; x$. The probability that less than X units are procured, i.e., the corresponding distribution in the fixed-quantity auction is

$$\tilde{G}(X) = \text{Prob.}(A \leq 2\, B^B(X/2)) = \frac{1}{12}\, X^2 + \frac{\sqrt{60}}{15}\, X.^9$$

The left panel of Figure 4.1 plots the difference between the bids in the two auction designs, i.e., $b^{FQ}(x) - b^B(x)$. Bids converge in the quantity supplied. Intuitively, the budget effect, which explains the difference between bids in the first place, becomes less relevant with increasing quantities. Both margins and the weight attached to them decrease in the quantity supplied. Consequently, the

[8] The convenient property that the relevant part of the bidding function spans the same interval also explains the choice of the budget auction as the baseline for the comparison.

[9] Bids in the corresponding fixed-quantity auction are $b^{FQ}(x) = \frac{1}{2}\,x - \frac{\sqrt{15}}{5} + \frac{27\sqrt{3}\arccos\left(\sqrt{15}/45(5x+2\sqrt{15})\right)}{2\sqrt{225-75x^2-60\sqrt{15}x}}$.

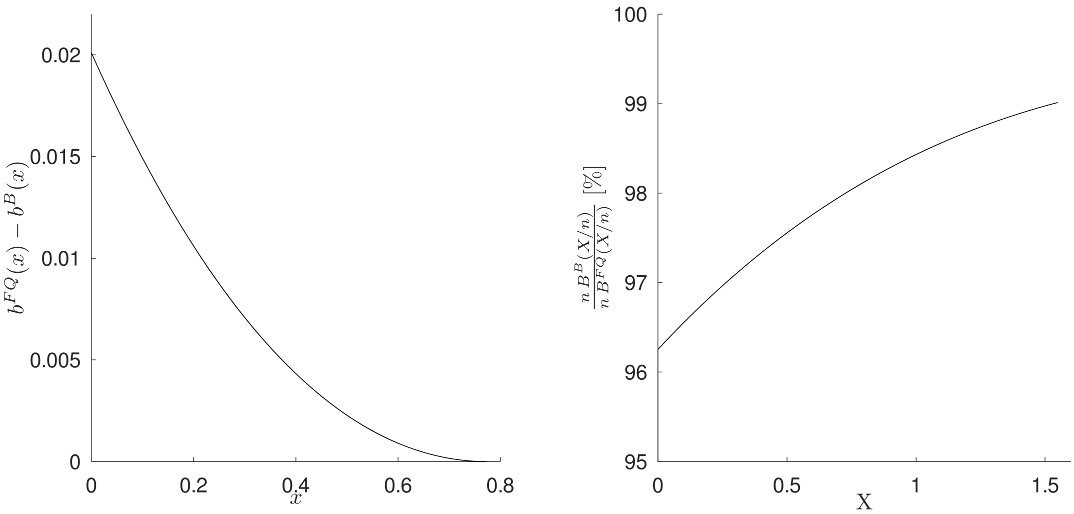

Figure 4.1: The difference between the equilibrium bids in the fixed-quantity auction and the equilibrium bids in the budget auction to which it corresponds ($n = 2$, $c(x) = x$ and $A \sim U[0, 1]$). On the left: difference in bids, on the right: total costs of procurement in the budget auction as a share of the total costs in the corresponding fixed-quantity auction.

budget effect gives less and less reason to lower bids to increase the probability of realising the margins on higher units as the quantity increases.

The right panel of Figure 4.1 shows the auctioneer's total cost of procurement in the budget auction as a share of the total procurement costs in the corresponding fixed-quantity auction, i.e., $\frac{n\,B^B(X/n)}{n\,B^{FQ}(X/n)}$. In this example, the procurement costs in the budget auction are between roughly 96% and 99% of the procurement costs in the fixed-quantity auction.

Evidently, the size of the cost advantage of the budget auction varies with the model parameters. In particular, we are interested in the impact that some specific model parameter, say n, has on the relation of total procurement costs under the two auction designs; i.e., we would like to know $\frac{d}{dn}\frac{B^B(x)}{B^{FQ}(x)}$. This requires gauging the relative size of $\frac{d}{dn}B^B(x)$ and $\frac{d}{dn}B^{FQ}(x)$ for $\forall x < \bar{x}$. However, the established approaches to doing analytic comparative "statics" on the equilibrium paths do not allow to do so. Even though Oniki (1973) generalises comparative statics via implicit differentiation of the first-order condition to the dynamic context, his approach only allows *signing* the relevant expressions by drawing a phase diagram of the corresponding system of linear differential equations. Similarly, the primal-dual methodology, comprehensively set out in Caputo (2005), uses the information contained in the second-order conditions to *sign* the effect of parameter changes on the optimal paths *integrated over the entire planning horizon*, i.e., $\int_0^{\bar{x}} \frac{d}{dn} B^{B/FQ}(x)\, dx$ in the case at hand. Thus, these approaches do not permit extracting the necessary information on the relative size of the impact of a parameter between the two auction designs in the absence of a closed-form solution. Consequently, the following only provides suggestive evidence on the economic contexts favourable to the budget auction based on numeric calculation.

The budget auction reduces competition problems faced by fixed-quantity auctions. Put differently, the cost advantage of the budget auction should be particularly large when fixed-quantity auctions suffer from soft competition. The general intuition is that bids in the budget auction are reduced to increase the probability of realising margins on other units. As a consequence, bids are reduced most strongly when the margins on these other units are high, i.e., competition is soft. Specifically, the budget auction should have a stronger cost advantage over the fixed-quantity auction when there are few bidders whose cost structure exhibits strong diseconomies of scale.

Figure 4.2 illustrates these points based on numerical calculation. The findings are robust to various specifications regarding the form of costs and the underlying distribution. To numerically construct a measure of the budget auction's cost advantage, we first need to determine equilibrium bids in the budget auction. To do this, we can guess an arbitrary value for $\bar{x}$, which fixes $B(\bar{x})$, $b(\bar{x})$, $b'(\bar{x})$ and $b''(\bar{x})$ by equations (4.6), (4.7), (4.8), and (4.A.8), respectively. Using the Euler equation from (4.5), we can develop the solution "backward," i.e., calculate bids

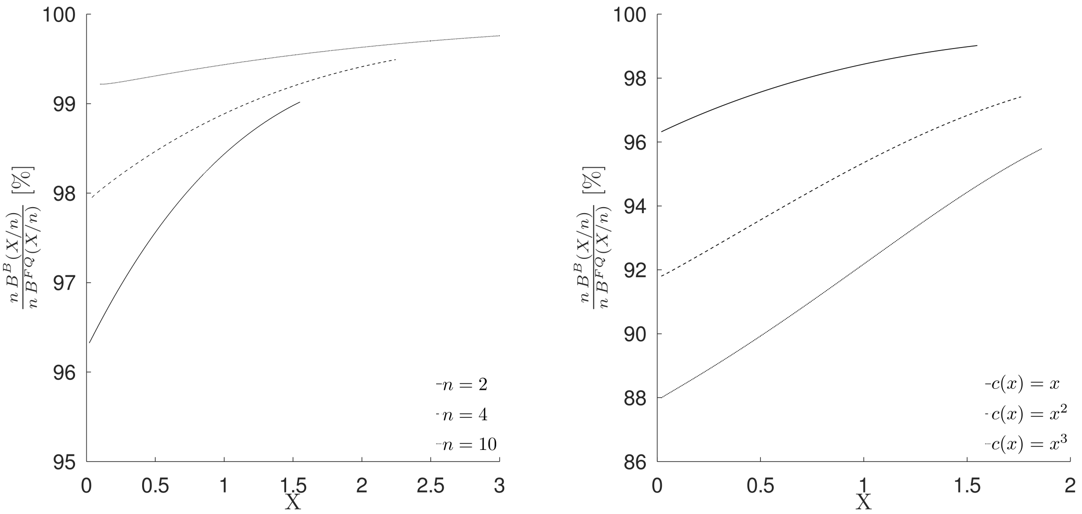

Figure 4.2: Total costs of procurement in the budget auction as a share of total costs in the corresponding fixed-quantity auction. On the left: comparison for a budget auction with $c(x) = x$, $A \sim U[0, 1]$, and a differing number of players. On the right, the comparison is shown for a budget auction with $n = 2$, $A \sim U[0, 1]$, and differing cost structures.

for ever lower quantities. The correct choice of $\bar{x}$ is identified by forcing $B(0) = 0$. We can then numerically determine the distribution $\tilde{G}(X)$, and equilibrium bids in the corresponding fixed-quantity auction.

The left panel of Figure 4.2 plots the cost advantage, i.e., $\frac{n\,B^B(X/n)}{n\,B^{FQ}(X/n)}$, for the previous example not only for the case of two (solid line) but also four and ten bidders. As can be seen, the fewer bidders there are, the lower the number of units the auctioneer can procure with a given budget.[10] Under both auction designs, a decreasing number of players softens competition and thus increases margins. Increased margins mean that the budget effect becomes more relevant. Therefore, the relative cost advantage of the budget auction increases when there are fewer bidders.

For the same reason, stronger diseconomies of scale increase the cost advantage of the budget auction, as can be seen in the right panel of Figure 4.2.[11] Again,

[10]This is due to increased competition and total social costs for producing a given quantity decreasing with more producers.

[11]Stronger diseconomies of scale here are interpreted as "more convex" cost functions. In the example, this is formalised analogously to utility functions with constant relative risk aversion.

the solid line corresponds to the previous example, and solutions are obtained numerically. To understand why diseconomies of scale increase profit margins, first consider the case of a flat marginal cost curve. In this case, the considered auctions would collapse into Bertrand competition with bids competed down to cost. It is only when marginal costs are increasing that bidders can realise profits because other bidders who already supply a larger quantity cannot exert full competitive pressure, as they are in a less favourable cost position. Extending this logic, the faster costs increase, i.e., the larger the diseconomies of scale, the higher the profit margins bidders can realise. As margins increase, the budget effect becomes more relevant, and the cost advantage of budget auctions increases.

4.5 Limitations

Uncertainty of the Budget.— There are two reasons for introducing uncertainty regarding the budget. From a theoretical point of view, without uncertainty, the equilibrium from Proposition 4.2 only persists under the assumption that the bidders bid below their costs on units they do not supply in equilibrium. As this is weakly dominated by an otherwise unchanged strategy with above-cost bids for off-equilibrium quantities, such an equilibrium does not seem plausible in real auctions. As the distribution converges to a point mass on $\bar{A}$, bids converge to flat bids at level $c(\bar{x})$. The reason is that there is no cost of making the own bids on the infra-marginal units the bidder might lose, less attractive, as the density attached to these units is zero. However, the bidders have an incentive to deviate upward in this case. Consider the incentive of a bidder i to marginally raise their constant bids from $b_i = c(\bar{x}_i)$. Their profits for a constant $b_i \geq c(\bar{x}_i)$ are $\pi_i = A - \sum_{j \neq i} B_j(\beta_j(b_i)) - C(\frac{A - \sum_{j \neq i} B_j(\beta_j(b_i))}{b_i})$, where $A - \sum_{j \neq i} B_j(\beta_j(c(\bar{x}_i))) \equiv A_i$.

The finding regarding bidders' cost is not driven by the more convex cost functions resulting in lower costs for relevant quantities.

If the other bidders bid weakly above costs for $x > \bar{x}_i$, the deviating bidder's profits are differentiable for a marginal upward deviation in b_i. Then, the benefit of marginally deviating upward is

$$\frac{\partial \pi_i}{\partial b_i}\bigg|_{b_i = c(A_i/b_i)} = \frac{\partial A_i}{\partial b_i} - c\left(\frac{A_i}{b_i}\right)\left(\frac{\partial A_i}{\partial b_i}\frac{1}{b_i} - \frac{A_i}{(b_i)^2}\right) = \frac{A_i}{b_i} > 0.$$

Intuitively, the bidder only loses a non-profitable marginal unit but gains a margin on all other units. It is easy to verify that the same holds for fixed-quantity auctions.

From a practical point of view, some literature suggests that clear focal points make collusive behaviour more likely (see, for example, Cramton and Ockenfels, 2017; Knittel and Stango, 2003). This is often a major concern of practitioners. An unknown budget can prevent collusion by eliminating the clear focal point of evenly splitting the budget. However, if procurement agencies are generally unable or unwilling to create uncertainty over the budget, the results of this chapter remain primarily theoretical.

Uncertainty of the Bidders' Cost.— The importance of an uncertain budget for the model hinges on the assumption of common knowledge of costs among the bidders. If their costs are private information, a symmetric equilibrium in pure strategies likely persists with a known budget, and collusion is harder as well. Consequently, an uncertain budget may not be necessary. Not considering private costs is the main limitation of the model in this chapter.

The budget effect probably carries over to a model with private costs and a known budget: A bidder raising their bids curbs the total demanded quantity and, thus, the quantity the bidder can win. Fundamentally, the bidders' problems share the same structure in both models, whether the probability of winning a certain number of units corresponds to the auctioneer's budget being sufficiently high or the other bidders not being too competitive.

Optimal Mechanism with Individual Bidder Types.— A model with individual bidder types and a known budget also allows for the analysis of the optimal mechanism for procuring the highest possible quantity with a given budget in a natural setup. Such a perspective complements the findings of this chapter. First, this chapter is not concerned with the *optimal* mechanism but with the strategic differences between budget and fixed-quantity auctions. Second, it does not model the auctioneer's uncertainty over the bidders' costs explicitly. Thus, it does not allow meaningfully answering the question of optimality.

If there is a continuum of bidder types, with high types having unambiguously higher costs than low types, and it is common knowledge that types represent independent draws from some distribution, the optimal mechanism for maximising the procured quantity with a given budget can be derived analogously to Dasgupta and Spulber (1989). Participation constraints and incentive compatibility have standard implications, with no information rents accruing to the least efficient type and information rents reflecting the cost advantage of more efficient types. However, the auctioneer always spends the same amount, i.e., the entire budget, as this allows them to procure a higher quantity. Consequently, when all bidders are of the least efficient type, demand is not distorted: as would be the case under perfect information, the auctioneer only pays for production costs and spends the entire budget. Thus, the same quantity as under perfect information is procured. By contrast, the bidders always make profits on inframarginal units in the budget auction.

Endogenous Quality.— The model also abstracts from quality considerations. These could weaken the result of Proposition 4.4 if the more competitive environment of the budget auction entailed a stronger incentive for bidders to lower product quality and costs. However, the marginal increase in expected profits a bidder can realise by lowering their costs is identical in both auction designs due to the dynamic equivalent of the standard envelope theorem (see Caputo, 1990).

As profits are maximised with respect to bids, a decrease in costs does not affect profits through bids. It only increases profits to the degree that it reduces costs for the (unchanged) equilibrium quantity. Thus, with the normalisation of demand in Proposition 4.4, reducing costs is equally attractive under both designs.

Preferences of the Auctioneer.— Lower unit costs in the budget auction do not imply that the auctioneer necessarily prefers the budget auction. The paper of Weitzman (1974) gives insights into how the preferences of the auctioneer, which are not modeled explicitly in this chapter, matter. In his model, the state either fixes a level of emissions or their price. The market outcome regarding the other variable is unknown to the state when it decides on its policy, as the costs of market participants are uncertain. In the model of Weitzman (1974), market participants do not behave strategically. Consequently, the costs of reaching a given level of emissions are identical under both policy options. Nonetheless, the state normally prefers one option: By fixing the level of emissions or their price for all realisations of costs, the state generally misses the optimal outcome in any given situation. Roughly speaking, the state should fix the price if the optimal price of emissions is almost identical for all cost realization, whereas the optimal level of emissions varies strongly and vice versa. An auctioneer who chooses between a budget and a fixed-quantity auction faces a similar problem. They fix either total spending or the traded quantity, and the not-fixed variable varies with the bidders' unknown costs. The insights of Weitzman (1974) imply that an auctioneer who wants to spend roughly the same amount (procure the same quantity) for any realisation of the bidders' cost by tendency prefers a budget (fixed-quantity) auction. For example, an auctioneer who intrinsically values spending exactly their budget (for example, because of internal budgeting) tends to prefer a budget auction. In contrast to the model of Weitzman (1974), an auctioneer choosing between a budget and a fixed-quantity auction must also account for the strategic differences between the two. As the budget auction allows procuring a given quantity at lower total costs, it gains in attractiveness compared to a setting without strategic effects for most plausible preferences of the auctioneer.

4.6 Conclusion

This chapter has considered a pay-as-bid procurement auction, in which the auctioneer maximises the number of procured units given a secret budget constraint. In line with previous literature, the analysis assumed that the bidders are symmetric and share a common belief concerning the distribution of the auctioneer's budget. Deriving the symmetric equilibrium of the budget auction revealed that the budget auction introduces additional strategic complexity in comparison to the case of an uncertain but fixed quantity. In the budget auction, bidders not only weigh higher profit margins against a lower probability of realising these same margins; higher bids also negatively influence the probability of winning later units. If a bidder demands a higher payment for some quantity, this reduces the auctioneer's residual budget such that it is less likely that the auctioneer's budget suffices to procure higher quantities from the bidder, which lowers bids. Thus, the budget cap leads to a link between bids in the bidders' problem, which has no correspondence in auctions with a fixed traded quantity. Due to this effect, bids in the budget auction are lower than in an auction with uncertain quantity, in which bidders believe in the same distribution of procured units as in the equilibrium of the budget auction.

4.A Appendix: Proof of Proposition 4.2

4.A.1 First-Order Condition

The relevant derivatives for the first-order condition given in (4.4) are

$$\frac{\partial \Psi}{\partial B_i} = -[b_i(x) - c(x)]\, f\Big(B_i(x) + \sum_{j \neq i} B_j\big(\beta_j(b_i(x))\big)\Big), \tag{4.A.1}$$

$$\begin{aligned}\frac{\partial \Psi}{\partial b_i} = {} & 1 - F\Big(B_i(x) + \sum_{j \neq i} B_j\big(\beta_j(b_i(x))\big)\Big) \\ & - [b_i(x) - c(x)]\, f\Big(B_i(x) + \sum_{j \neq i} B_j\big(\beta_j(b_i(x))\big)\Big) \sum_{j \neq i} b_i(x)\, \beta_j'(b_i(x)),\end{aligned} \tag{4.A.2}$$

and

$$\begin{aligned}\frac{d}{dx}\left(\frac{\partial \Psi}{\partial b_i}\right) = {} & -f(\cdot)\, \Big[b_i(x) + \sum_{j \neq i} b_i(x)\, \beta_j'(b_i(x))\, b_i'(x)\Big] \\ & - \Big[b_i'(x) - c'(x)\Big]\, f(\cdot)\, \sum_{j \neq i} b_i(x)\, \beta_j'(b_i(x)) \\ & - \Big[b_i(x) - c(x)\Big] \Bigg\{ f'(\cdot)\, \Big[b_i(x) + \sum_{j \neq i} b_i(x)\, \beta_j'(b_i(x))\, b_i'(x)\Big] \\ & \times \sum_{j \neq i} b_i(x)\, \beta_j'(b_i(x)) + f(\cdot)\, \sum_{j \neq i} b_i'(x)\, \beta_j'(b_i(x)) \\ & + f(\cdot)\, \sum_{j \neq i} b_i(x)\, \beta_j''(b_i(x))\, b_i'(x) \Bigg\}.\end{aligned} \tag{4.A.3}$$

Searching for a symmetric equilibrium with $b(x) = b_i(x) = b_j(x)\, \forall j$, furthermore gives $\beta_j'\big(b_i(x)\big) = \frac{1}{b'(x)}$ and $\beta_j''\big(b_i(x)\big) = -\frac{b''(x)}{[b'(x)]^3}$. Using these symmetry conditions together with the derivatives in equations (4.A.1) and (4.A.3) in the necessary condition from equation (4.4) gives equation (4.5).

4.A.2 Definitising the (numeric) Solution to the Characteristic Differential Equation

Definitising $B(\bar{x})$ and $b(\bar{x})$. — The highest quantity a given bidder can win, is determined according to the terminal curve $B_i(\bar{x}_i) = \phi(b_i, \bar{x}_i)$. Based on the general transversality condition for the endpoint of a variational problem, the Euler-Lagrange equation of the considered bidder's best response problem from equation (4.4) is completed by $B_i(\bar{x}_i) = \phi(b_i, \bar{x}_i)$ and $[\Psi(\cdot) + (\frac{d}{d\bar{x}_i}\phi(\cdot) - b_i(x))\frac{\partial\Psi}{\partial b_i}]_{x=\bar{x}_i} = 0$ (see e.g., Chiang, 1992). These conditions translate to

$$B_i(\bar{x}_i) + \sum_{j \neq i} B_j\Big(\beta_j(b_i(\bar{x}_i))\Big) = \bar{A} \tag{4.A.4}$$

and given equation (4.A.4)

$$\Big(\sum_{j \neq i} b_i(\bar{x}_i)\, \beta_j'(b_i(\bar{x}_i))\, b_i'(\bar{x}_i) + b_i(x)\Big)\ [b_i(\bar{x}_i) - c(\bar{x}_i)]\, f(\bar{A}) \sum_{j \neq i} b_i(\bar{x}_i)\, \beta_j'(b_i(x)) = 0. \tag{4.A.5}$$

By assumption $f(\bar{A}) > 0$ and bids of all players are positive and increasing, such that equation (4.A.5) implies that

$$b_i(\bar{x}_i) = c(\bar{x}_i). \tag{4.A.6}$$

For the symmetric equilibrium, the transversality conditions from equations (4.A.4) and (4.A.6) give equations (4.6) and (4.7).

Definitising $b'(\bar{x})$.— Given that for a response of i to be optimal, it must be that $b_i(\bar{x}_i) = c(\bar{x}_i)$, the first-order condition of bidder i's problem obtained by using equations (4.A.1) and (4.A.3) in equation (4.4) simplifies substantially. Making use of the assumptions that $f(\bar{A})$, $f'(\bar{A})$, as well as the derivatives of the bidding functions, are finite, several terms vanish such that the simplified first-order condition emerges as

$$b_i'(\bar{x}_i) = \frac{c'(\bar{x}_i)}{2} - \frac{1}{2\ \sum_{j \neq i} \beta_j'\Big(c(\bar{x}_i)\Big)}. \tag{4.A.7}$$

As, in equilibrium, this condition must hold for every player, we can use symmetry to obtain the formulation in equation (4.8).

Definitising $b''(\bar{x})$.— The formulation of $b''(\bar{x})$, as given in (4.5), cannot be used for the numerical approximation of the symmetric equilibrium, it is indeterminate. This can be seen by substituting the expression for $b'(\bar{x})$ from (4.8) and noting that $b(\bar{x}) - c(\bar{x}) = 0$. Applying Hôpital's rule to (4.5) yields

$$b''(\bar{x}) = \frac{[b'(\bar{x})]^2}{b(\bar{x})\,[b'(\bar{x}) - 2c'(\bar{x})]}\Bigg[3b'(\bar{x}) - \frac{(2n-3)}{(n-1)}\,c'(\bar{x}) - c''(\bar{x})\,\frac{b(\bar{x})}{b'(\bar{x})} + n\,\frac{f'[nB(\bar{x})]}{f(nB(\bar{x}))}\,[b'(\bar{x}) - c'(\bar{x})]\,\frac{[b(\bar{x})]^2}{b'(\bar{x})}\Bigg]. \tag{4.A.8}$$

4.A.3 Properties of Equilibrium Bids

Bids are Above Marginal Costs.— Consider a bidder i with bidding function $b_i(x)$ who, on some arbitrarily small, relevant interval $[x_1; x_2]$ of strictly positive length, bids weakly below marginal costs. Furthermore, assume that the bidder also has a weakly positive margin on the interval $[x_2; x_3]$ with $x_3 - x_2 = x_2 - x_1$ and $x_3 \le \bar{x}_i$. Consider the alternative bidding function

$$\hat{b}_i(x) = \begin{cases} b_i(x) + \varepsilon & \text{for } x \in [x_1; x_2] \\ b_i(x) - \delta\varepsilon & \text{for } x \in (x_2; x_3] \\ b_i(x) & \text{else,} \end{cases}$$

where $\delta \ge 1$ and $x_3 = [(1+\delta)x_2 - x_1)]/\delta$. For an appropriate choice of δ the bidder would profit from choosing a positive ε. Because both $b_i(x)$ and $B_i(x)$

are unchanged outside the interval $[x_1; x_3]$, the marginal effect on profits of increasing ε from 0 is

$$\begin{aligned}\frac{\partial \mathbb{E}[\pi_i]}{\partial \varepsilon}\bigg|_{\varepsilon=0} = &\int_{x_1}^{x_2} \left[1 - F\Big(B_i(x) + \sum_{j\neq i} B_j\big(\beta_j(b_i(x))\big)\Big)\right] dx - \int_{x_2}^{x_3(\delta)} \delta[1-F(\cdot)]\, dx \\ &- \underbrace{\int_{x_1}^{x_2} [b_i(x) - c(x)]\, f(\cdot) \left[(x - x_1) + \sum_{j\neq i} b_i(x)\beta_j'(b_i(x))\right] dx}_{\leq 0} \\ &+ \int_{x_2}^{x_3(\delta)} [b_i(x) - c(x)]\, f(\cdot)\, \delta \sum_{j\neq i} b_i(x)\beta_j'(b_i(x))\, dx \\ &- \int_{x_2}^{x_3(\delta)} [b_i(x) - c(x)]\, f(\cdot)\, [(x_2 - x_1) - \delta(x - x_2)]\, dx. \end{aligned} \tag{4.A.9}$$

For high enough $\delta, \frac{\partial \mathbb{E}[\pi_i]}{\partial \varepsilon}|_{\varepsilon=0} > 0$. To see this, note that the first three lines are each (weakly) positive. The first line is positive because $\int_{x_1}^{x_2} 1 - F(\cdot)dx > (x_2 - x_1)[1-F(\cdot)]_{x=x_2} = \delta(x_3 - x_2)[1-F(\cdot)]_{x=x_2} > \int_{x_2}^{x_3} \delta[1-F(\cdot)]dx$. The second line is positive due to the weakly negative margins in the interval. Lastly, all terms in the integrand in the third line are positive. The integral in the fourth line is decreasing in δ and converges to zero as $\delta \to \infty$. It follows that for high enough δ, the expression in (4.A.9) is positive and there is a profitable deviation from $b_i(x)$.[12] This argument also holds if the bidder does not bid above costs for $x > x_2$. As a consequence, the profitability of a deviation to $\hat{b}_i(x)$ shows that a bidder's best response cannot include below-cost bidding.

Monotonicity of Bids.— Equation (4.8) establishes that $b'(\bar{x}) > 0$. Against the backdrop of equation (4.B.2) the slope of equilibrium bids for $x \in [0, \bar{x})$ can be written as

$$b'(x) = b'_{FQ}(x) - \Delta b'(x).^{13} \tag{4.A.10}$$

[12] The strategy $\hat{b}_i(x)$ is of course inadmissible as a solution because it has jump discontinuities at x_2 and x_3. However, the *optimal solution* is continuous, as can easily be seen from setting the bidder's problem up as a problem of optimal control, which would allow jump discontinuities in $b_i(x)$.

[13] Note that this proof is out of proper logical order.

Given equations (4.B.1) and (4.B.3), we can write

$$b'(x) = (n-1)\,[1-F(\cdot)]^{-1}\,f(\cdot)\,b(x)\,[b(x)-c(x)] - (n-1)\,[1-F(\cdot)]^{\frac{1-n}{n}}\,\Delta\tilde{b}'(x) > 0. \tag{4.A.11}$$

It was shown above that $b(x) > c(x)$ for $x \in [0,\bar{x})$ and Appendix 4.B shows that $\Delta\tilde{b}'(x) < 0$ for $x \in [0,\bar{x})$. Thus, equation (4.A.11) establishes the monotonicity of bids.

4.A.4 Second-Order Condition

General Case.— The most straightforward way to ensure that the symmetric equilibrium candidate described by equations (4.5) - (4.8) does indeed constitute an equilibrium is to make sure that the integrand of a unilaterally deviating bidder's profit function is globally concave. Although, depending on the case at hand, weaker conditions might be given, I will state the conditions for the strict concavity of $\Psi(\cdot)$. For the deviating bidder i, the integrand of their expected profits is

$$\Psi(B_i, b_i, x) = [b_i(x) - c(x)]\left[1 - F\Big(B_i(x) + (n-1)\,B\big(\beta(b_i(x))\big)\Big)\right]. \tag{4.A.12}$$

Strict global concavity requires that for any B_i and b_i the corresponding quadratic form is negative definite. In particular, it must be the case that

$$\frac{\partial^2 \Psi}{\partial (b_i)^2} < 0 \tag{4.A.13}$$

and

$$\frac{\partial^2 \Psi}{\partial (b_i)^2} \times \frac{\partial^2 \Psi}{\partial (B_i)^2} - \left(\frac{\partial^2 \Psi}{\partial b_i\,\partial B_i}\right)^2 < 0 \tag{4.A.14}$$

for all (B_i, b_i), where

$$\begin{aligned}\frac{\partial^2 \Psi}{\partial (b_i)^2} &= -2(n-1)\,f(\cdot)\,b_i(x)\,\beta'(\cdot) - [b_i(x)-c(x)]\;f'(\cdot)\;\Big[(n-1)\,b_i(x)\,\beta'(\cdot)\Big]^2 \\ &\qquad - (n-1)\,[b_i(x)-c(x)]\,f(\cdot)\,\Big(\beta'(\cdot) + b_i(x)\,\beta''(\cdot)\Big),\\ \frac{\partial^2 \Psi}{\partial (B_i)^2} &= -[b_i(x)-c(x)]\;f'(\cdot)\end{aligned}$$

and

$$\frac{\partial^2 \Psi}{\partial b_i\, \partial B_i} = -f(\cdot) - (n-1)\,[b_i(x) - c(x)]\, f'(\cdot)\, b_i(x)\, \beta'(\cdot).$$

The condition of global concavity can be weakened, as any maximum of a bidder's expected profits can only be located in the area where $b_i(x) > c(x)$.

Example from Section 4.3.3.— Assume that the continuation of the bidding function for off-equilibrium quantities is differentiable and weakly concave. The problem satisfies concavity given that only cases with $b_i(x) - c(x)$ need to be considered. In this case $\frac{\partial^2 \Psi}{\partial (b_i)^2} < 0$. Furthermore, due to the uniform distribution $\frac{\partial^2 \Psi}{\partial (B_i)^2} = 0$ and $\frac{\partial^2 \Psi}{\partial b_i\, \partial B_i} = -1$. Thus, the second-order condition is satisfied in the example.

4.B Appendix: Proof of Proposition 4.4

The main text establishes that $b^{FQ}(\bar{x}) = b^B(\bar{x})$. It remains to be shown that $b^{FQ}(x) > b^B(x)$ for $x \in [0;\bar{x})$. Based on equations (4.3.2) and (4.9) bids in the corresponding fixed-quantity auction are

$$b^{FQ}(x) = (n-1)\Big[1 - F\Big(n\, B^B(x)\Big)\Big]^{\frac{1-n}{n}} \int_x^{\bar{x}} [1 - F(\cdot)]^{-1/n} f(\cdot)\, b^B(y)\, c(y)\, dy. \quad (4.B.1)$$

Noting that bids in the budget auction can be written as

$$b^B(x) = (n-1)\,\Big[1 - F\Big(n\, B^B(x)\Big)\Big]^{\frac{1-n}{n}} \mu(x) \int_x^{\bar{x}} [1 - F(\cdot)]^{-1/n} f(\cdot)\, b^B(y)\, b^B(y)\, dy,$$

where

$$\mu(x) = \frac{b^B(x)}{b^B(x) + \Big[1 - F\Big(nB^B(x)\Big)\Big]^{\frac{1-n}{n}} \int_x^{\bar{x}} \Big[1 - F\Big(nB^B(y)\Big)\Big]^{\frac{n-1}{n}} b^{B\prime}(y)\, dy},$$

the difference in bids, $\Delta b(x) = b^{FQ}(x) - b^B(x)$, is

$$\Delta b(x) = (n-1)\,[1 - F(\cdot)]^{\frac{1-n}{n}} \int_x^{\bar{x}} \Big[1 - F(\cdot)\Big]^{-1/n} f(\cdot)\, b^B(y)\, [c(y) - \mu(x)\, b^B(y)]\, dy. \quad (4.B.2)$$

Given that we are considering $x < \bar{x}$, the hypothesis that $\Delta b(x) > 0$ is equivalent to

$$\Delta\tilde{b}(x) = \int_x^{\bar{x}} \Big[1 - F\Big(n\,B^B(y)\Big)\Big]^{-1/n} f\Big(nB^B(y)\Big)\, b^B(y)\, [c(y) - \mu(x)\, b^B(y)]\, dy > 0. \tag{4.B.3}$$

Because $\lim_{x\to\bar{x}} \Delta\tilde{b}(x) = 0$, a sufficient condition for inequality (4.B.3) and by extension $\Delta b(x) > 0$ for $x < \bar{x}$ to hold is that $\Delta\tilde{b}'(x) < 0$, i.e.,

$$\begin{aligned}\Delta\tilde{b}'(x) = [1 - F(\cdot)]^{-1/n} f(\cdot)\, b^B(x)\Big[\mu(x)\, b^B(x) - c(x)\Big] \\ - \mu'(x) \int_x^{\bar{x}} [1 - F(\cdot)]^{-1/n} f(\cdot)\, [b^B(y)]^2\, dy < 0.\end{aligned} \tag{4.B.4}$$

Dividing by $[1 - F(\cdot)]^{-1/n}\; f(\cdot)\, b^B(x)$, integrating the second line by parts, and plugging in the derivative of $\mu(x)$ reveals that inequality (4.B.4) is equivalent to

$$b^B(x) - c(x) - \frac{1}{(n-1)}\, \frac{1 - F\Big(nB^B(x)\Big)}{f\Big(nB^B(x)\Big)}\, \frac{b^{B\prime}(x)}{b^B(x)} < 0. \tag{4.B.5}$$

A comparison with equation (4.A.2) shows that inequality (4.B.5) is equivalent to $\frac{\partial\Psi}{\partial b_i}\Big|_{B_i=B_j=B^B, b_i=b_j=b^B} > 0$. Given that bids are above costs (see Appendix 4.A.3), equation (4.A.1) reveals that $\frac{\partial\Psi}{\partial B_i} < 0$ for $x \in [0, \bar{x})$ if evaluated along the equilibrium path. Thus, the first-order condition from (4.4) implies that $\frac{d}{dx}\frac{\partial\Psi}{\partial b_i} < 0$ for $x < \bar{x}$ in equilibrium. Using the terminal conditions from equations (4.6) and (4.7) in (4.A.2) gives that, evaluated at equilibrium bids, $\frac{\partial\Psi}{\partial b_i}\Big|_{x=\bar{x}} = 0$. It follows that $\frac{\partial\Psi}{\partial b_i}\Big|_{B_i=B_j=B^B, b_i=b_j=b^B} > 0$, i.e., inequality (4.B.5) holds.

5

Watt's Better? Pay-as-Bid vs. Uniform Pricing in Electricity Markets

A version of this chapter (see Appendix B) has been invited for resubmission at Energy Economics.

Abstract

Rising electricity prices of recent years have reopened the debate on replacing uniform with pay-as-bid pricing in electricity markets. This chapter contributes to this ongoing political debate by comparing peak prices, consumer surplus and welfare under pay-as-bid and uniform pricing in a theoretical model. For this comparison, it derives supply function equilibria for both pricing rules in a model that matches four stylised facts of electricity markets: Suppliers have (1) oligopolistic market power, (2) increasing marginal costs, (3) face a downward-sloping demand, and (4) have uncertainty over time-varying demand but common knowledge of production costs. In the model, peak prices are lower under pay-as-bid pricing. Both pay-as-bid and uniform pricing achieve first-best welfare with zero profits when marginal costs are flat, or there is an infinite number of producers. Restricting attention to the case where marginal costs and demand are linear with a uniformly distributed intercept of demand, pay-as-bid pricing results in a higher expected consumer surplus. The welfare comparison is ambiguous even in this linear model. Pay-as-bid pricing results in higher expected welfare if and only if demand uncertainty is sufficiently low. The findings of this chapter suggest that regulators should seriously consider pay-as-bid pricing to raise consumer surplus and mitigate price peaks.

5.1 Introduction

On January 8th, 2025, the price of a MWh of electricity on the intraday market for Great Britain peaked at more than £1,000/MWh (Ambrose, 2025) – roughly ten times the going rate for the whole month (EPEX, 2025), and several times the estimated production costs for even the most expensive technologies (Department for Energy Security & Net Zero, 2023a; International Energy Agency, 2020). Such price peaks raise questions about the adequacy of the current market institutions, particularly the way prices are set in electricity markets.

Currently, under the uniform-price rule, every unit trades at the same price, which is determined by the price of the most expensive unit. Intuitively, this approach places the burden of high-price situations disproportionately on consumers, who pay high prices even for units that are inexpensive to produce. Consequently, this price rule has come under fire amid the rising energy prices of recent years. For example, in her State of the Union address in 2022, Ursula von der Leyen, president of the European Commission, argued that "[t]he current electricity market design – based on merit order – is not doing justice to consumers anymore. They should reap the benefits of low-cost renewables. So, we have to decouple the dominant influence of gas [usually the most expensive technology] on the price of electricity" (European Commission, 2022). An obvious way to do so would be to change the pricing rule to pay-as-bid pricing, where every unit trades at "its own" price, which corresponds to the ask-price of the producer. Thus, regulators have recently started debating whether pay-as-bid pricing should replace uniform pricing.[1]

[1] For example, the UK government's Review of Electricity Market Arrangements considered and rejected such a switch because of "the risk of tactical bidding"(Department for Energy Security & Net Zero, 2023b). This chapter compares price rules accounting for this strategic behaviour.

This chapter contributes to this ongoing debate by comparing peak prices, consumer surplus and welfare under pay-as-bid and uniform pricing in a theoretical model.

The key challenge in the formal analysis of firms' equilibrium behaviour in electricity markets is that they submit bids as combinations of quantities and corresponding prices, rather than choosing a single price or quantity. Thus, a proper analysis of electricity markets requires finding so-called "supply function equilibria". Doing so is particularly complex for pay-as-bid pricing, where a firm's revenue depends on its entire supply function, as each unit is remunerated at the bid price for that unit. So far, the literature has not established equilibrium behaviour under assumptions that plausibly reflect electricity markets.

This chapter is the first work to derive equilibrium bids under pay-as-bid pricing in a model that incorporates four important stylised facts of electricity markets in a single model. First, there are typically only a few suppliers. Thus, electricity suppliers strategically exert market power (e.g., Borenstein *et al.*, 2002; Bushnell *et al.*, 2008; Krzywnicka and Barner, 2025, among others), as opposed to the models in Federico and Rahman (2003), Song *et al.* (2025), Willems and Yueting (2023), or Zhao *et al.* (2023). Second, even though demand is rather inelastic, it does react to prices even in the very short term (Hirth *et al.*, 2024), which, for example, the models in Genc (2008), Hästö and Holmberg (2006), Holmberg (2008, 2009), Pycia and Woodward (2026), and Zhao *et al.* (2023) do not reflect. Third, suppliers are well-informed about their competitors' costs/bidding behaviour (Doraszelski *et al.*, 2018) but face uncertainty over the exact level of demand (see Laitsos *et al.*, 2024, for a review of the literature on forecasting demand in electricity markets). Some literature analyses minimal information environments regarding competitors' cost (e.g., Galgana and Golrezaei, 2025; Kasberger and Woodward, 2025) or abstracts from uncertainty over demand (e.g., Aussel *et al.*, 2017a,b; Caragiannis *et al.*, 2025;

Galgana and Golrezaei, 2025; Kasberger and Woodward, 2025; Son *et al.*, 2004; Vanelli *et al.*, 2025).[2] Fourth, some models (e.g., Caragiannis *et al.*, 2025; Fabra and Llobet, 2023; Fabra *et al.*, 2006; Genc, 2008; Holmberg and Wolak, 2018) do not account for the fact that – apart from purely renewable electricity suppliers – suppliers usually have increasing marginal costs (their so-called "merit order") as they own a portfolio of power plants with differing production technologies (see, e.g., Borenstein and Bushnell, 1999; Hortaçsu and Puller, 2008, for corresponding empirical estimates).

Accommodating these stylised facts in a single model, this chapter derives the supply function equilibrium under pay-as-bid pricing when suppliers with commonly known, (weakly) increasing marginal costs compete for a time-varying downward-sloping demand in an oligopolistic market. For uniform pricing, the equilibrium is known from Klemperer and Meyer (1989). Comparing pay-as-bid and uniform pricing in this general setup, this chapter finds that peak prices are lower under pay-as-bid pricing. Moreover, both pricing rules implement first-best welfare with zero profits under perfect competition, i.e., if marginal costs are flat or there is an infinite number of competitors. To derive insights into consumer surplus and welfare outside these extremely competitive cases, the chapter also considers a linear version of the model, where marginal costs and demand are linear with a uniformly distributed intercept of demand. In this linear model, the expected consumer surplus is higher under pay-as-bid pricing. Expected welfare increases through pay-as-bid pricing if and only if the uncertainty over demand is sufficiently low.

This chapter proceeds as follows. Section 5.2 discusses related literature. Section 5.3 presents the general model and derives equilibrium bids for arbitrary demand and cost functions. Based on this equilibrium characterisation, it shows

[2]Kasberger and Woodward (2025) and Vanelli *et al.* (2025) do so because they do not consider electricity markets.

that peak prices are lower under pay-as-bid pricing. Both price rules implement the first-best under perfect competition. Section 5.4 compares consumer surplus and welfare in the linear model. Section 5.5 discusses the implications of the model, before Section 5.6 concludes.

5.2 Related Literature

Finding optimal bidding strategies in electricity markets is mathematically complex. Many studies rely on simulation models where computer agents learn through repeated interactions (e.g., Bakirtzis and Tellidou, 2006; Bower and Bunn, 2001; Guerci *et al.*, 2007; Hailu and Thoyer, 2007; Liu *et al.*, 2012; Sugianto and Liao, 2014; Viehmann *et al.*, 2021; Xiong *et al.*, 2004). A strength of this approach is that it can incorporate market features such as grid constraints (Guerci and Rastegar, 2012) or dynamic technological learning (Anatolitis and Welisch, 2017), which are difficult to capture in analytic models. On the downside, such studies only identify equilibrium behaviour in a specific setting with numeric values for relevant parameters such as costs and demand. Thus, it remains unclear under which conditions their findings generalise.

Analytical models partly address this concern by deriving equilibrium behaviour mathematically, making the role of specific assumptions more transparent. In the model of this chapter, *oligopolistic* electricity suppliers with *increasing* marginal costs compete for a *time-varying/uncertain*, *downward-sloping* demand. Hence, the model expands on previous literature that has not considered these aspects in a single model.[3] Each of these assumptions is relevant for electricity markets and has important implications for equilibrium behaviour.

[3] Outside the context of electricity markets, some authors also restrict attention to single-unit demand/supply of bidders (e.g., Anderson and Holmberg, 2023; Bougt *et al.*, 2025; Krishna, 2010).

A minority of the related literature compares uniform and pay-as-bid pricing under the assumption that electricity suppliers are price-takers and do not strategically influence prices (e.g., Federico and Rahman, 2003; Song *et al.*, 2025; Willems and Yueting, 2023; Zhao *et al.*, 2023). Although this is certainly true for a subset of suppliers, the empirical evidence shows that models of imperfect competition better describe behaviour in electricity markets (e.g., Borenstein *et al.*, 2002; Bushnell *et al.*, 2008; Krzywnicka and Barner, 2025, among others). As is intuitively plausible, both pricing rules perform identically and maximise welfare under perfect competition in the model of this chapter (Corollary 5.2).

Many other models assume that demand is completely inelastic (e.g., Genc, 2008; Holmberg, 2008, 2009; Hästö and Holmberg, 2006; Pycia and Woodward, 2026; Zhao *et al.*, 2023). Even though electricity demand really is rather inelastic, it does react to price increases even in the (very) short term (see Hirth *et al.*, 2024, for a recent analysis, or Labandeira *et al.*, 2017, for a meta-analysis). Accounting for this demand-side reaction is important as it limits the exercise of market power (Cramton, 2004). For example, Heim and Götz (2021) empirically document very high prices under pay-as-bid pricing in the German market for reserve power, where regulation implies that demand is completely inelastic. Similarly, models of uniform-price auctions with inelastic demand regularly feature *binding* price caps (e.g., Genc, 2008; Holmberg, 2008, 2009; Zhao *et al.*, 2023).[4] The model of this chapter nests the case of inelastic demand as a limiting case.

Some models assume that marginal production costs are constant (e.g., Caragiannis *et al.*, 2025; Fabra and Llobet, 2023; Fabra *et al.*, 2006; Genc, 2008; Holmberg and Wolak, 2018). Generally, this assumption does not accurately reflect that suppliers usually own a fleet of power plants with differing production technologies

[4]In the model of this chapter, inelastic could lead to infinite prices under uniform pricing but not pay-as-bid pricing.

and, therefore, have increasing marginal costs (see, e.g., Borenstein and Bushnell, 1999; Hortaçsu and Puller, 2008, for corresponding empirical estimates). When comparing consumer surplus under pay-as-bid and uniform pricing, accounting for increasing marginal costs is crucial, as the main disadvantage of uniform pricing for consumers is that they pay high prices for low-cost units when a high-cost technology sets the price. By contrast, the assumption of constant (zero) marginal costs becomes plausible for markets approaching 100% renewable electricity generation. Such purely renewable markets introduce a new complexity because suppliers have uncertainty over their own production capacity, which, for example, depends on meteorological conditions (see Zhao *et al.*, 2023, for such a model). If renewable electricity suppliers can eliminate uncertainty over their own production capacity through aggregation of production facilities (e.g., Song *et al.*, 2025), such 100% renewable markets correspond to the limiting case of this model with constant marginal costs. However, if this uncertainty persists, such markets deserve an entirely separate treatment, as uncertainty over their own production introduces fundamentally new considerations into bidding behaviour (Fabra and Llobet, 2023).

The last important component of electricity market models is the information structure. Some recent literature (e.g., Galgana and Golrezaei, 2025; Kasberger and Woodward, 2025) considers minimal information environments. Bidders operate under "deep uncertainty" over competitors' costs and cannot even attach probabilities to different possibilities. Thus, they do not maximise expected profits. Instead, they minimise their loss in profits compared to a situation where they know their competitors' costs (regret). However, such models are mostly adequate for very rare interactions (Kasberger and Woodward, 2025). As suppliers interact repeatedly in electricity markets, they are well-informed about their competitors' costs (Hortaçsu and Puller, 2008), especially as all producers employ similar production technologies. Thus, empirically, suppliers in electricity seem to bid according

	Oligopoly	Elastic Demand	Increasing Marginal Costs	Stochastic Uncertainty
Aussel *et al.* (2017a,b)	✓		✓	
Caragiannis *et al.* (2025)	✓			
Fabra *et al.* (2006)	✓	(✓) (extension)		(✓) over demand (extension)
Federico and Rahman (2003)	(✓) only monopoly	✓	(✓) only across producers	✓ over demand
Galgana and Golrezaei (2025)	✓		✓	deep uncertainty on comp.bids
Genc (2008)	✓			✓ over demand
Hästö and Holmberg (2006), Holmberg (2009)	✓		✓	✓ over demand
Holmberg and Wolak (2018)	✓			✓ over own and competitors' cost and demand
Kasberger and Woodward (2025)	✓		✓	deep uncertainty on comp. bids
Pycia and Woodward (2026)	✓		✓	✓ over demand
Son *et al.* (2004)	✓		✓	
Song *et al.* (2025)		random prices		✓ over equilibrium prices
Vanelli *et al.* (2025)	✓	✓	✓	
Willems and Yueting (2023)		✓	(✓) only across producers	✓ over demand
Zhao *et al.* (2023)				✓ over own production capacity
This chapter	✓	✓	✓	✓ over demand

Table 5.1: Overview of assumptions in the related literature.

to a Nash equilibrium, where they know their competitors' costs (Doraszelski *et al.*, 2018).[5] In line with the substantial literature on demand forecasting (and its error) (see Laitsos *et al.*, 2024, for a review) and much of the theoretical literature (e.g., Federico and Rahman, 2003; Genc, 2008; Holmberg, 2009; Holmberg and Wolak, 2018; Hästö and Holmberg, 2006; Willems and Yueting, 2023), this chapter assumes that suppliers are left with uncertainty over the time-varying demand. Some models (e.g., Aussel *et al.*, 2017a,b; Caragiannis *et al.*, 2025; Son *et al.*, 2004; Vanelli *et al.*, 2025) also eliminate this source of uncertainty. Incorporating demand uncertainty not only seems empirically more plausible but also leaves the no-uncertainty case as the limiting case where the distribution of demand approaches a point mass.

Table 5.1 presents the assumptions made in previous literature. As discussed, the novelty of this chapter lies in its integration of the plausible electricity market

[5] The close alignment of cost curves estimated from bidding data with engineering estimates in Wolak (2003) also indirectly supports this view, suggesting that externally accessible engineering estimates of costs can reliably predict bids.

assumptions of supply-side market power, increasing marginal costs, and time-varying, downward-sloping demand. According to Table 5.1, this chapter is most closely related to the work of Federico and Rahman (2003), who compare pay-as-bid and uniform pricing with linear marginal costs and random, linear demand in a monopoly. This chapter extends their work in two directions: First, Section 5.3 derives equilibrium behaviour in oligopoly without any of these functional form restrictions. Second, Section 5.4 extends their comparison with linear functional forms to the setting of oligopolistic market power.

5.3 General Model

Demand.— The demanded quantity in the entire market at price p is given by $D(p,\varepsilon)$, where ε is a demand-shifter. Without loss of generality, assume that $\frac{\partial D(p,\varepsilon)}{\partial \varepsilon} > 0$. The demand-shifter ε is a random variable that follows the cumulative distribution $F(\varepsilon)$ on $[\underline{\varepsilon}, \bar{\varepsilon}]$ with $F'(\varepsilon) = f(\varepsilon) > 0$. This chapter assumes that demand is downward-sloping, i.e., that $\frac{\partial D(p,\varepsilon)}{\partial p} < 0$. However, the case of inelastic demand is nested in the derivation of equilibrium behaviour. Throughout the chapter, assume that $D(p,\varepsilon)$ reflects consumers' preferences without strategic distortion.[6]

Supply.— There are $n \geq 2$ symmetric producers that compete in the market. Each producer has the same, commonly known marginal cost function $c(x)$ with $c'(x) \geq 0$.[7] Every bidder i submits a weakly increasing bidding function $\beta_i(x)$ that reflects, how much money they demand when supplying x units.

Market Clearing.— The market clears at a stop-out price p^* that equalises demand and aggregate supply, i.e.,

$$D(p^*, \varepsilon) = \sum_i \beta_i^{-1}(p^*), \tag{5.1}$$

[6]Thus, the model abstracts from the strategic reaction to pricing rules on the demand side (see Kamat and Oren, 2002, for a model with strategic demand).

[7]Equilibrium bidding strategies do not change when including fixed costs.

where $\beta_i^{-1}(p)$ is the inverse of $\beta_i(x)$. Thus, as long as bidding functions are strictly increasing, a producer supplies $\beta_i^{-1}(p^*)$ units, i.e., all units for which their bidding function is weakly below p^*. Throughout the chapter, assume that the minimal demand level is high enough for trade to occur for every demand realisation under both pricing rules.

Pricing Rules.— Under the uniform-price rule, a buyer is paid the stop-out price p^* for all units they supply. When they bid $\beta_i(x)$ and the stop-out price is p^*, they make profits of

$$\pi_{UP} = p^* \, \beta_i^{-1}(p^*) - C\Big(\beta_i^{-1}(p^*)\Big),$$

where $C(x) = \int_0^x c(y)\,dy$. Under the pay-as-bid rule, they are paid the price specified by their bidding function $\beta_i(x)$ for each supplied unit; their profits with stop-out price p^* are

$$\pi_{PAB} = \int_0^{\beta_i^{-1}(p^*)} \beta_i(x)\,dx - C\Big(\beta_i^{-1}(p^*)\Big).$$

Figure 5.1 illustrates a producer's revenue under the two pricing rules.

5.3.1 Symmetric Equilibrium under Pay-as-Bid Pricing

Bidder's Problem.— Under pay-as-bid pricing, a producer i earns their margin $\beta_i(x) - c(x)$ on some unit x if and only if they get to supply weakly more than x units. Assume all other producers $j \neq i$ follow the strategy $\beta_j(x)$. In that case, producer i supplies weakly more than x units, if and only if demand at price $\beta_i(x)$ is large enough to cover the x units of producer i and the $(n-1)\,\beta_j^{-1}(\beta_i(x))$ units the other producers supply at this price. Thus, producer i earns the margin $\beta_i(x) - c(x)$ if and only if $D(\beta_i(x), \varepsilon) \geq x + (n-1)\beta_j^{-1}(\beta_i(x))$. This is equivalent to the demand-shifter ε taking a high enough value, specifically,

$$\varepsilon \geq e\Big(\beta_i(x), x + (n-1)\beta_j^{-1}(\beta_i(x))\Big),$$

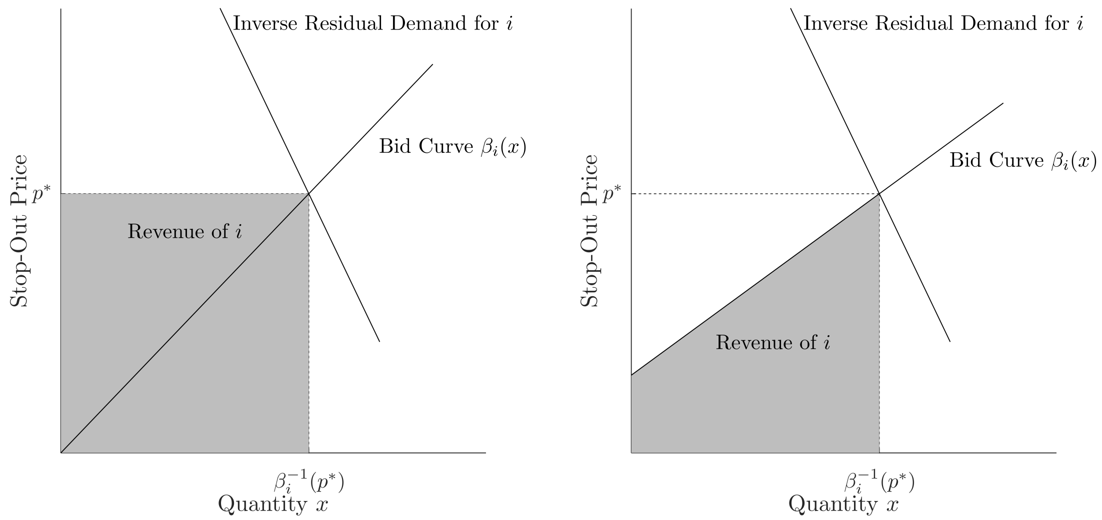

(a) Uniform pricing. **(b)** Pay-as-bid pricing.

Figure 5.1: Illustration of uniform and pay-as-bid pricing.

where $e(p, Q)$ it the unique value of ε for which $D(p, e) = Q$. Thus, bidder i's equilibrium bidding strategy solves the maximisation problem

$$\max_{\beta_i(x)} \qquad \int_{\underline{x}_i}^{\bar{x}_i} [\beta_i(x) - c(x)] \left[1 - F\Big(e\Big(\beta_i(x), x + (n-1)\beta_j^{-1}(\beta_i(x))\Big)\Big)\right] dx, \quad (5.2)$$

where $\underline{x}_i$ and $\bar{x}_i$ denote the quantities bidder i gets to supply with their strategy when demand is at its minimal level $\underline{\varepsilon}$ and maximum level $\bar{\varepsilon}$, respectively. The objective function in (5.2) corresponds to the margin on unit x, weighted by the probability that the demand-shifter ε is large enough for the producer to supply more than x units.

The maximisation problem in (5.2) can be solved by point-wise maximisation of the integrand; the first-order condition is

$$1 - F(\cdot) = [\beta_i(x) - c(x)]\, f(\cdot)\Big(e_p(\cdot) + e_Q(\cdot)(n-1)\beta_j^{-1\prime}(\beta_i(x))\Big), \qquad (5.3)$$

where $e_p(\cdot)$ and $e_Q(\cdot)$ denote the partial derivatives of $e(\cdot)$ with respect to p and Q, respectively. The producer benefits from increasing their bid on some unit x by making additional profits whenever unit x is inframarginal, i.e., whenever they win more than x units (left-hand side); on the downside, they lose the margin on the unit in the case where they do not get to supply the unit because of the higher bid, i.e., when the unit is marginal (right-hand side). As the unit $\bar{x}_i$, which the producer only supplies when $\varepsilon = \bar{\varepsilon}$, is never inframarginal, there is no incentive to bid above marginal costs on that unit. Consequently, producer i's optimal bidding function satisfies

$$\beta_i(\bar{x}_i) = c(\bar{x}_i).^8 \tag{5.4}$$

The trade-off in the first-order condition from (5.3) only captures the producer's problem correctly, as long as they are at risk of losing the unit in question: In that case, the producer trades off a positive probability of realising the margin on that unit (positive left-hand side) with the positive probability of losing that margin if demand is sufficiently low (positive right-hand side). However, for units $x < \underline{x}_i$, there is no risk of losing them, as the producer supplies $\underline{x}_i$ units even if demand is at its minimal level. Evidently, the supplier should increase the bid on these units until they become marginal for the minimal demand level. Thus, the bid on all units $x \leq \underline{x}_i$ is identical and the optimal bidding schedule of producer i is

$$\beta_i(x) = \begin{cases} b_i(\underline{x}_i) & \text{for } x < \underline{x}_i \\ b_i(x) & \text{for } x \geq \underline{x}_i, \end{cases} \tag{5.5}$$

where $b_i(x)$ denotes the solution to (5.3) and $\underline{x}_i$ is the quantity they supply at the minimal demand level under this solution, i.e., $\underline{x}_i = D\Big(b_i(\underline{x}_i), \underline{\varepsilon}\Big) - (n - 1)\beta_j^{-1}\Big(b_i(\underline{x}_i)\Big)$.

Bidding Behaviour in the Symmetric Equilibrium.— Equations (5.3)–(5.5) describe the optimal bid function of an individual bidder i, whose competitors all follow the same bidding strategy $\beta_j(x)$. Based on this description of optimal bids,

[8] Formally, the left-hand side of (5.3) is zero for $\varepsilon = \bar{\varepsilon}$. Thus, the right-hand side must be zero to satisfy (5.3). Given that with increasing bids all but the first factor are positive, (5.4) must hold.

Proposition 5.1 characterises the symmetric equilibrium, i.e., the situation in which bidder i's best response it to follow the strategy $\beta_j(x)$ as well.

Proposition 5.1 (Equilibrium under Pay-as-Bid Pricing) *Consider a pay-as-bid auction with $n \geq 2$ symmetric bidders with commonly known, weakly increasing marginal costs $c(x)$. Assume that market demand is given by some, weakly downward-sloping demand function $D(p, \varepsilon)$ such that $D_\varepsilon(\cdot) > 0$ where ε follows the cumulative distribution $F(\cdot)$ on $[\underline{\varepsilon}, \bar{\varepsilon}]$. The symmetric equilibrium in pure strategies is given by*

$$\beta^*_{PAB}(x) = \begin{cases} b^*_{PAB}(\underline{x}_{PAB}) & \text{for } x < \underline{x}_{PAB} \\ b^*_{PAB}(x) & \text{for } x \geq \underline{x}_{PAB}, \end{cases} \tag{5.6a}$$

with

$$b^{*\prime}_{PAB}(x) = \frac{(n-1)\, e_Q\Big(b^*_{PAB}(x), nx\Big)\, [b^*_{PAB}(x) - c(x)]\, f\Big(e(\cdot)\Big)}{1 - F\Big(e\Big(b^*_{PAB}(x), nx\Big)\Big) - e_p(\cdot)\, [b^*_{PAB}(x) - c(x)]\, f(\cdot)}, \tag{5.6b}$$

$$b^*_{PAB}(\bar{x}_{PAB}) = c(\bar{x}_{PAB}), \tag{5.6c}$$

$$\bar{x}_{PAB} = \frac{D(c(\bar{x}_{PAB}), \bar{\varepsilon})}{n}, \tag{5.6d}$$

and

$$\underline{x}_{PAB} = \frac{D(b^*_{PAB}(\underline{x}_{PAB}), \underline{\varepsilon})}{n}, \tag{5.6e}$$

where $e(p, Q)$ is such that $D\Big(p, e(p, Q)\Big) = Q$.

Proof. Equations (5.6a)–(5.6c) follow from using the definition of a symmetric equilibrium – i.e., $x_i = x_j = x$, $p = \beta(x) = \beta_i(x) = \beta_j(x)$ and, consequently, $\beta^{-1\prime}(p) = 1/\beta'(x)$ – in equations (5.3)–(5.5). Equations (5.6d) and (5.6e) formalise the definition of the minimal and maximum quantity a bidder receives in the symmetric equilibrium. ■

If demand is completely inelastic, we have that $e(p, Q) = Q$, in which case (5.6b) corresponds to the solution in Pycia and Woodward (2026).

5.3.2 Symmetric Equilibrium under Uniform Pricing

This section briefly revisits equilibrium behaviour under uniform-pricing, which has been established by Klemperer and Meyer (1989). They show that – for given bids of the other suppliers – a bidder's equilibrium bids maximise the bidder's profits for every realisation of the demand-shifter ε, i.e., they are an ex-post equilibrium.

Bidder's Problem.— Therefore, we can model an individual producer i as maximising their profits for a given realisation of ε. If the other producers $j \neq i$ bid according to $\beta_j(p)$, producer i will supply $x_i = D(p,\varepsilon) - (n-1)\beta_j^{-1}(p)$ units at price p. Thus, they solve

$$\max_p \qquad \pi_{UP} = \Big(D(p,\varepsilon) - (n-1)\beta_j^{-1}(p)\Big)\,p - C\Big(D(p,\varepsilon) - (n-1)\beta_j^{-1}(p)\Big). \quad (5.7)$$

The first-order-condition of this problem is

$$x_i = \Big(p - c(x_i)\Big)\,\Big((n-1)\beta_j^{-1\prime}(p) - D_p(p,\varepsilon)\Big), \quad (5.8)$$

where $D_p(\cdot)$ is the derivative of $D(\cdot)$ with respect to p. The left-hand side of (5.8) reflects the upside of marginally raising the price, or equivalently, raising the bid on the x_i-th unit; the producer obtains additional marginal earnings on the x_i inframarginal units. However, they also lose marginal units and the corresponding margins, which the right-hand side of (5.8) reflects. Hence, producers bid truthfully on the first unit, as they do not benefit from increased margins on any inframarginal units. Proposition 5.2 describes bids in a symmetric equilibrium under uniform pricing based on the first-order condition in (5.8).

Proposition 5.2 (Klemperer and Meyer, 1989: Equilibrium under Uniform Pricing) *The bidding strategies in a symmetric equilibrium of a uniform-price auction with $n \geq 2$ symmetric bidders with commonly known, weakly increasing marginal* *costs $c(x)$, when market demand is given by some demand function $D(p, \varepsilon)$ satisfy*

$$\beta^{*\prime}_{UP}(x) = \frac{(n-1)[\beta^*_{UP}(x) - c(x)]}{x + D_p\Big(\beta^*_{UP}(x), \varepsilon\Big)[\beta^*_{UP}(x) - c(x)]} \tag{5.9a}$$

$$\beta^*_{UP}(0) = c(0). \tag{5.9b}$$

Proof. Use the definition of a symmetric equilibrium, i.e., $x_i = x_j = x$, $p = \beta(x) = \beta_i(x) = \beta_j(x)$, and $\beta^{-1\prime}(p) = 1/\beta'(x)$ in the first-order condition from (5.8). ■

Multiplicity of Equilibria and Equilibrium Selection.— Klemperer and Meyer (1989) show that there are infinitely many solutions to the initial value problem in (5.9), i.e., there are infinitely many symmetric equilibria under uniform pricing. Intuitively, producers' pay-offs only depend on a single point on their bid function for a given demand realisation; the rest of the bidding function stabilises the outcome (see Back and Zender, 2001).

When comparing price rules, this leaves a problem of equilibrium selection. Klemperer and Meyer (1989) show that there is (potentially) only one equilibrium, in which producers – for sufficiently high demand – (i) offer their entire, infinite production capacity and (ii) do so at weakly positive margins. The rest of this chapter focuses on such equilibria. All other equilibria seem – at least to some extent – implausible in electricity markets. If (ii) is violated, margins become negative at some point, resulting in negative profits for high demand levels.[9] If an equilibrium violates (i), there is a strict upper bound to the number of units the capacity-unconstrained producers supply; for high enough demand levels, producers would even curb the quantity they supply as demand increases.[10] Competition authorities most likely would prevent such behaviour. Pycia and Woodward (2026) interpret the remaining equilbria as uncertainty-robust equilibria; no matter, how much

[9]Of course, real-world producers do not have infinite production capacity. Thus, some of the equilibria that violate (ii) might be plausible in the real world.

[10]Formally, equilibria violating (i) and (ii) correspond to the denominator and the numerator in (5.9a) going to zero such that the slope of bids goes to zero or infinity, respectively.

demand potentially increases through the demand-shifter ε, i.e., independently of the distribution $F(\varepsilon)$, these equilibria remain economically meaningful. This robustness makes these equilibria focal for the comparison to pay-as-bid pricing.

5.3.3 Peak Prices, Consumer Surplus, and Welfare

Peak Prices.— One of the objectives of switching to pay-as-bid pricing is to avoid price peaks, such as the one in January 2025, which was described in Section 5.1. The equilibrium behaviour under pay-as-bid and uniform pricing established in Propositions 5.1 and 5.2 directly implies that peak prices are lower under pay-as-bid pricing.

Intuitively, peak prices, i.e., prices when demand is at its maximum, depend on the bidders' margins for large quantities. Under pay-as-bid pricing, bidders do not make a positive margin on the price-setting unit when demand is at its maximum. They only benefit from a higher margin on that unit when demand is at its maximum, which – with a continuous distribution of the demand shock – has a probability weight of zero. By contrast, uniform pricing results in particularly high margins in that situation: Raising bids raises the margin on the large number of inframarginal units.

Corollary 5.1 formalises this intuition.

Corollary 5.1 (Lower Peak Prices under Pay-as-Bid Pricing) *Consider an auction with $n \geq 2$ symmetric bidders with commonly known, weakly increasing marginal costs $c(x)$. Assume that market demand is given by some, weakly downward-sloping demand function $D(p, \varepsilon)$ such that $D_\varepsilon(\cdot) > 0$ where ε follows the cumulative distribution $F(\cdot)$ on $[\underline{\varepsilon}, \bar{\varepsilon}]$. If profitable trade is possible, peak prices, i.e., prices when demand is $D(p, \bar{\varepsilon})$, are strictly lower under pay-as-bid pricing than in the uncertainty-robust equilibrium under uniform price auction.*

Proof. Equilibrium bidding functions are weakly increasing (for uniform pricing: in an uncertainty-robust equilibrium). Thus, the highest prices result for $\varepsilon = \bar{\varepsilon}$. The quantity an individual bidder supplies in that case solves $D\Big(\beta^*_{PAB/UP}(x_{PAB/UP}), \bar{\varepsilon}\Big) = n\,\bar{x}_{PAB/UP}$. By (5.6c), $\beta^*_{PAB}(\bar{x}_{PAB}) = c(\bar{x}_{PAB})$. If profitable trade is possible, $\bar{x}_{PAB} > 0$. Corollary 5.1 claims that $\beta^*_{UP}(\bar{x}_{UP}) > \beta^*_{PAB}(\bar{x}_{PAB})$. Prove this inequality by contradiction. Suppose that $\beta^*_{UP}(\bar{x}_{UP}) \leq \beta^*_{PAB}(\bar{x}_{PAB})$. If peak prices were weakly lower under uniform pricing, it would need to be the case that $\bar{x}_{UP} \geq \bar{x}_{PAB}$ because demand is (weakly) downward-sloping. As costs are weakly increasing, it follows that $c(\bar{x}_{UP}) \geq c(\bar{x}_{PAB}) = \beta^*_{PAB}(\bar{x}_{PAB})$. Thus, peak prices are only lower under uniform pricing if $\beta^*_{UP}(\bar{x}_{UP}) \leq c(\bar{x}_{UP})$, which – with increasing bids – is inconsistent with optimal bidding behaviour under uniform pricing for $\bar{x}_{UP} > 0$ according to the first-order condition in (5.8). ■

While peak prices play a focal role in the public debate, they are not an accurate measure of the performance of a price rule. They only focus on an extreme demand situation and, under pay-as-bid pricing, do not even measure actual market prices for this demand situation, as inframarginal units trade at a lower price. To address these limitations, the following considers consumer surplus and welfare as comprehensive measures of the performance of price rules.

Consumer Surplus and Welfare under Perfect Competition.— The minimal functional form assumptions of the model impose significant constraints on ranking the two pricing rules regarding consumer surplus and welfare. In fact, Section 5.4 shows that the welfare ranking is ambiguous even under the substantially more restrictive functional form assumptions of the linear model. However, Corollary 5.2 shows that both pay-as-bid and uniform pricing result in first-best welfare and zero profits under perfect competition, i.e., if there is an infinite number of suppliers or marginal costs are constant.

Corollary 5.2 (Welfare & Consumer Surplus under Perfect Competition) *Consider $n \geq 2$ symmetric producers with the commonly known marginal cost curve $c(x)$, where $c(x) \geq 0$ and $c'(x) \geq 0$. The producers compete in a uniform-price or a pay-as-bid auction with market demand at price p given by the downward-sloping demand function $D(p, \varepsilon)$, where ε is a random variable. For $n \to \infty$, or constant marginal costs, both auction formats implement the first-best welfare. Producers do not make profits and consumer surplus corresponds to first-best welfare.*

Proof. Consider the case of infinitely many suppliers. Under both pricing rules, the quantity supplied by an individual supplier, x_i, converges to zero as $n \to \infty$ for every demand realisation. It follows from Proposition 5.2 that the entire traded quantity is provided to consumers at price $\beta^*_{UP}(0) = c(0)$ under uniform pricing. Similarly, it follows from Proposition 5.1 that the entire traded quantity is provided at costs $\beta^*_{PAB}(\bar{x}_{PAB}) = c(\bar{x}_{PAB}) = c(0)$ to consumers under pay-as-bid pricing. The results on welfare and consumer surplus follow immediately.

Next, consider the case of constant unit costs. In that case, producers bid their true costs under both pricing rules. Thus, both pricing rules trade the first-best quantity and, given that marginal costs are constant, suppliers make zero profits. For pay-as-bid pricing, the solution to (5.6b) with constant marginal costs is $b^*_{PAB}(x) = c(x)$. Producers can only make positive margins to the degree that their competitors cannot undercut this price because they have higher costs when they produce more units. With constant marginal costs undercutting is always profitable and the pay-as-bid auction collapses into Bertrand competition. Similarly, it is easy to verify that truthful bidding is also an equilibrium under uniform pricing; it solves the initial value problem in (5.9). More precisely, truthful bidding is the uncertainty-robust equilibrium. ■

The result of Corollary 5.2 for constant marginal cost is the elastic-demand equivalent of the result in Genc (2008). By contrast, the result that pay-as-bid pricing achieves first-best welfare when there is an infinite number of suppliers contrasts with the findings of Federico and Rahman (2003), who conclude that the pay-as-bid auction does not implement first-best welfare under perfect competition. As noted by Willems and Yueting (2023), even though Federico and Rahman (2003) assume that suppliers are price-takers, their perfect competition benchmark reflects monopolistic competition rather than traditional perfect competition. In their model, every producer corresponds to a point on the increasing industry cost curve. Consequently, producers compete with competitors who have higher costs. Thus, less efficient competitors cannot undercut the price offered by a more efficient supplier even if that supplier has a positive margin. By contrast, in the model of this chapter, there is an infinite number of competitors, who can supply an additional unit at a price of $c(0)$. Hence, it is impossible to charge a higher price.

5.4 Linear Model

To derive more insights regarding consumer surplus and welfare outside the extremely competitive cases considered in Corollary 5.2, this section considers a linear model, where

$$D(p, \varepsilon) = \varepsilon - mp, \tag{5.10a}$$

$$\varepsilon \sim U[\underline{\varepsilon}, \bar{\varepsilon}], \tag{5.10b}$$

and

$$c(x) = \alpha x \tag{5.10c}$$

with $m, \alpha, \underline{\varepsilon} > 0$ and $\bar{\varepsilon} > \underline{\varepsilon}$.

5.4.1 Equilibrium Bids and Market Outcomes

Equilibrium Bids in the Linear Model.— Corollary 5.3 characterises the symmetric equilibria under pay-as-bid and uniform pricing in the linear model from (5.10).

Corollary 5.3 (Equilibrium Bids in the Linear Model) *Consider the case where demand is $D(p,\varepsilon) = \varepsilon - mp$ with $\varepsilon \sim U[\underline{\varepsilon},\bar{\varepsilon}]$ and there are $n \geq 2$ symmetric producers, each with the commonly known marginal cost curve $c(x) = \alpha x$ with $\alpha, m > 0$. The symmetric equilibrium in pure strategies in the pay-as-bid auction is given by the bidding function*

$$\beta^*_{PAB}(x) = \begin{cases} \lambda + \delta_{PAB}\,\underline{x}_{PAB} & \text{for } x < \underline{x}_{PAB} \\ \lambda + \delta_{PAB}\,x & \text{for } x \geq \underline{x}_{PAB} \end{cases} \tag{5.11a}$$

with

$$\lambda = \frac{(\alpha - \delta_{PAB})\bar{\varepsilon}}{(n + m\alpha)}, \tag{5.11b}$$

$$\delta_{PAB} = \frac{\sqrt{2\alpha m\,(2n-3) + (1-2n)^2 + \alpha^2 m^2} - 2n + \alpha m + 1}{4m}, \tag{5.11c}$$

and

$$\underline{x}_{PAB} = \frac{\underline{\varepsilon} - \lambda}{n + m\delta_{PAB}} \tag{5.11d}$$

assuming that the traded quantity is positive even with minimal demand, i.e., $\underline{\varepsilon} > \lambda$. Bids in the only symmetric equilibrium in pure strategies under uniform pricing, in which bids are increasing for all $x \in [0,\infty)$, i.e., in the uncertainty-robust equilibrium, satisfy

$$\beta^*_{UP}(x) = \delta_{UP}\,x = \frac{1}{2m}\Big(2 + \alpha\,m - n + \sqrt{4 + \alpha^2 m^2 - 4n + 2\alpha m n + n^2}\Big)\,x. \tag{5.12}$$

Proof. Noting that $e(p,Q) = Q + m\,p$ for the linear model from (5.10), it is straightforward to verify that the bid function in (5.11) satisfies the equilibrium conditions from Proposition 5.1. Similarly, it is straightforward to verify that (5.12) satisfies the initial value problem from Proposition 5.2 and is a uncertainty-robust equilibrium. The proof of its uniqueness is analogous to that for the linear example with two suppliers in Klemperer and Meyer (1989). ■

Figure 5.2 illustrates the equilibrium bids from Corollary 5.3. Bids in the linear model are linear. Under uniform pricing, bids start at marginal costs but then continue steeper than marginal costs.[11] For higher quantities, raising the price becomes increasingly attractive, as it increases the price on many inframarginal units. By contrast, under pay-as-bid pricing, suppliers have large margins on low quantities, with bids being less steep than costs. For the last unit, which the producer only supplies in the highest-demand situation, bids converge to marginal costs.[12]

Market Outcomes.— With equilibrium bids in place, we can characterise market outcomes. In the symmetric equilibrium, an individual producer gets to supply $D(\beta(x), \varepsilon) - (n-1)x$ units when the demand level is ε. Thus, they supply

$$x^*_{UP}(\varepsilon) = \frac{\varepsilon}{n + m\delta_{UP}}$$

and

$$x^*_{PAB}(\varepsilon) = \frac{\varepsilon - m\lambda}{n + m\delta_{PAB}}$$

units under uniform and pay-as-bid pricing, respectively. Consequently, the market trades $n\, x^*_{UP}(\varepsilon)$ and $n\, x^*_{PAB}(\varepsilon)$ at a stop-out price of $\beta^*_{UP}(x^*_{UP}(\varepsilon))$ and $\beta^*_{PAB}(x^*_{PAB}(\varepsilon))$, respectively. By contrast, under the first-best, when the traded quantity and stop-out price are determined by marginal costs, an individual producer would supply

$$x^*_{FB}(\varepsilon) = \frac{\varepsilon}{n + m\alpha}$$

units, and the market would trade $n\, x^*_{FB}(\varepsilon)$ units at a stop-out price of $\alpha\, x^*_{FB}(\varepsilon)$.

[11] Formally, bound the term in the square root in (5.12) by $(\alpha m + n - 2)^2$ and $(\alpha m + n)^2$, respectively, to get that $\alpha < \delta_{UP} < \alpha + 1/m$.

[12] Formally, replace the term in the square root in (5.11c) with $(\alpha m + 2n - 3)^2$ and $(\alpha m + 2n - 1)^2$, respectively, to verify that $(\alpha m - 1)/2m < \delta_{PAB} < \alpha/2$.

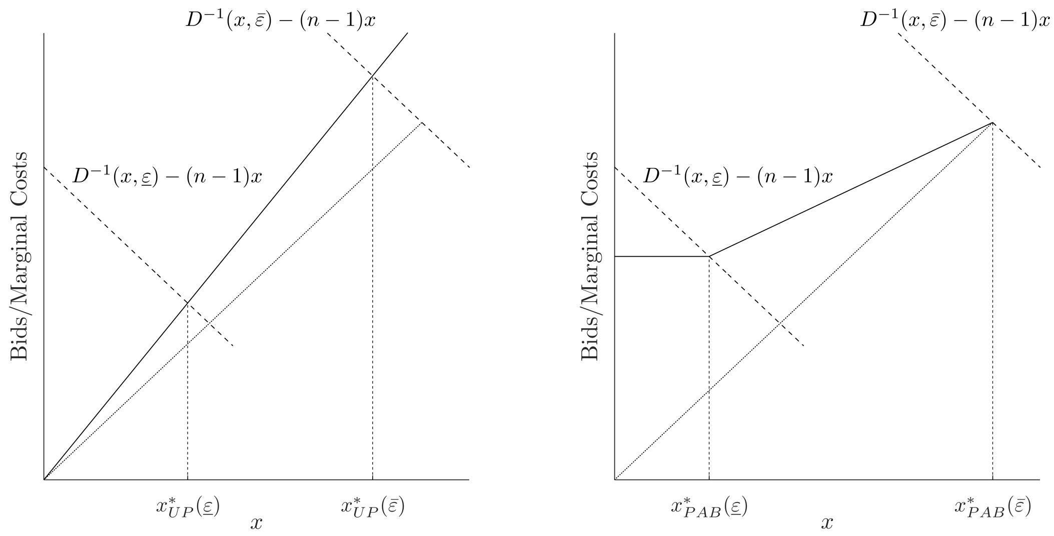

(a) Uniform pricing. **(b)** Pay-as-bid pricing.

Figure 5.2: Bids (solid line) and costs (dotted line) in the linear model under uniform and pay-as-bid pricing.

5.4.2 Consumer Surplus and Welfare

This characterisation of market outcomes makes it possible to compare consumer surplus and welfare in the linear model. Without loss of generality, the comparisons assume that $m = 1$.

Consumer Surplus.— Figure 5.3 illustrates consumer surplus for different demand levels under both pricing rules. It shows that the comparison of consumer surplus becomes more favourable to pay-as-bid pricing as demand increases: First, bids are steeper under uniform pricing, such that the stop-out price under uniform pricing increases faster in demand. Second, the rising stop-out price is more detrimental to consumer surplus under uniform pricing. In contrast to pay-as-bid pricing, consumers pay the increasing stop-out price on all units. Consequently, consumer surplus is higher under pay-as-bid pricing for the highest demand realisation. The stop-out price is lower under pay-as-bid pricing (see Corollary 5.1) and, additionally, consumers pay less than the stop-out price on inframarginal units. By contrast, they pay the higher stop-out price on all units under uniform

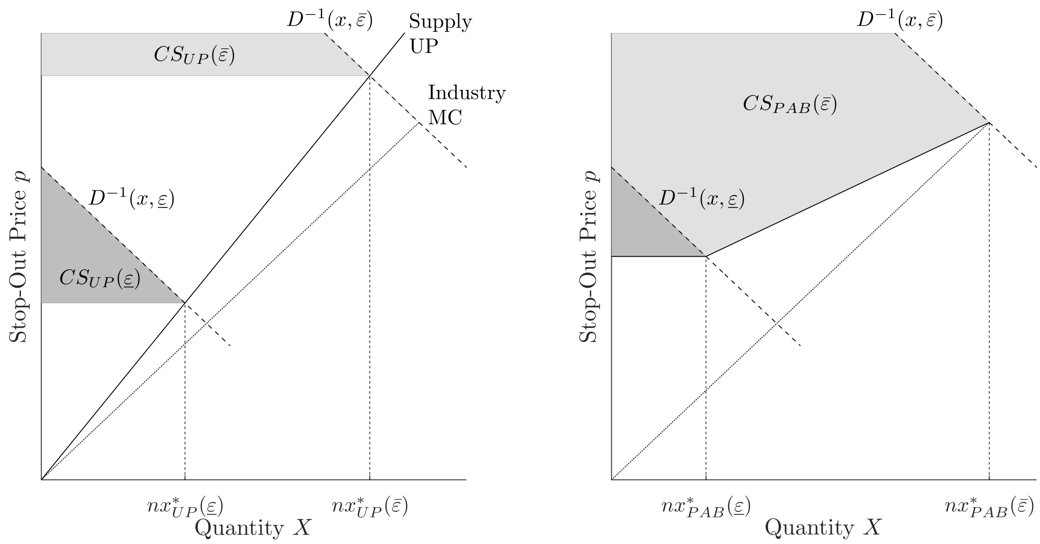

(a) Uniform pricing. **(b)** Pay-as-bid pricing.

Figure 5.3: Consumer surplus under uniform and pay-as-bid pricing. The dotted and solid lines represent industry-wide marginal costs and supply, respectively; the darkly and lightly shaded areas correspond to the consumer surplus for the lowest and highest demand realisation, respectively.

pricing. When demand is low, the comparison of consumer surplus is ambiguous. In Figure 5.3, uniform pricing realises a higher consumer surplus in the minimal demand situation but this is not necessarily the case.

A meaningful comparison of pricing rules must account for all demand situations, which requires comparing expected consumer surplus over all possible demand realisations. Formally, consumer surplus under uniform pricing for a given demand realistion is $CS_{UP}(\varepsilon) = \int_0^{nx^*_{UP}} D^{-1}(x,\varepsilon)\,dx - nx^*_{UP}\,\beta^*_{UP}(x^*_{UP})$, where $D^{-1}(x,\varepsilon)$ is the inverse demand function for a given demand level ε. Building the expectation over all demand realisations, the expected consumer surplus under uniform pricing is

$$\mathbb{E}[CS_{UP}] = \frac{n^2}{2(n+\delta_{UP})^2}\,\mathbb{E}[\varepsilon^2].$$

Under pay-as-bid pricing, consumer surplus for a given demand level is $CS_{PAB}(\varepsilon) = \int_0^{nx^*_{PAB}} D^{-1}(x,\varepsilon)\,dx - n\int_0^{x^*_{PAB}} \beta^*_{PAB}(x)\,dx$. In expectation, this gives a consumer surplus of

$$\mathbb{E}[CS_{PAB}] = \frac{n}{2(n+\delta_{PAB})}\,\mathbb{E}[\varepsilon^2] - \frac{n\lambda}{(n+\delta_{PAB})}\,\mathbb{E}[\varepsilon] + \frac{n\Big(n\lambda^2 - \underline{\varepsilon}(\underline{\varepsilon}-2\lambda)\delta_{PAB}\Big)}{2(n+\delta_{PAB})^2}.$$

Proposition 5.3 demonstrates that, in the expectation of all demand realisations, consumer surplus is higher under pay-as-bid than under uniform pricing, i.e., $\mathbb{E}[CS_{PAB}] > \mathbb{E}[CS_{UP}]$. As is intuitive from Figure 5.3, the proof shows that the reason is that consumers pay less than the stop-out price on inframarginal units.

Proposition 5.3 (Linear Model: Higher Expected Consumer Surplus under Pay-as-Bid-Pricing) *Consider the case where demand is* $D(p,\varepsilon) = \varepsilon - mp$ *with* $\varepsilon \sim U[\underline{\varepsilon}, \bar{\varepsilon}]$ *and there are* $n \geq 2$ *symmetric producers, each with the commonly known marginal cost curve* $c(x) = \alpha x$ *with* $\alpha, m > 0$. *Without loss of generality, assume that* $m = 1$. *Expected consumer surplus is strictly higher under pay-as-bid pricing than in the uncertainty-robust equilibrium of the uniform-price auction assuming that the traded quantity in the pay-as-bid auction is positive, i.e., that* $\underline{\varepsilon} \geq \lambda$.

Proof. See Appendix 5.A. ■

The result of Proposition 5.3 is unsurprising in the context of theoretical literature. For example, Pycia and Woodward (2026) show that for inelastic demand the *optimally designed* pay-as-bid auction results in a higher expected consumer surplus than the *optimally designed* uniform price auction without making any substantial functional form assumptions.[13] The big difference to the result in Proposition 5.3 is that they have the auctioneer choose the distribution of demand (and stochastic reserve prices) to maximise the expected consumer surplus under either pricing rule. Thus, their result does not say much about the policy choice between pricing rules in electricity markets: In the context of electricity markets, regulators have to take demand and its variation as given. Proposition 5.3 speaks to that situation.

[13]This is, translating their results to the reverse auction.

Welfare.— The comparison of expected welfare is generally ambiguous. Comparing welfare between the two pricing rules is equivalent to comparing the deadweight losses, i.e., the welfare loss compared to the first-best. Figure 5.4 represents the deadweight loss for varying demand levels under both pricing rules as shaded triangles. Formally, Figure 5.4 makes clear that the deadweight loss under pay-as-bid and uniform pricing is

$$DWL_{PAB/UP} = 0.5 \cdot (\beta^*_{PAB/UP}(x^*_{PAB/UP}) - \alpha\, x^*_{PAB/UP}) \cdot (n\, x^*_{PAB/UP} - n\, x_{FB}). \quad (5.13)$$

Equation (5.13) shows that welfare only depends on how close the traded quantity is to its first-best level (third factor).[14] The actual payment for electricity is only a welfare-neutral transfer from consumers to producers.

Under either pricing rule, the traded quantity converges to its first-best when the stop-out price is close to marginal costs. Thus, as illustrated in Figure 5.4, pay-as-bid pricing eliminates the deadweight loss when demand is at its maximum, as the bid on the "last" unit is truthful. More generally, the deadweight loss decreases with an increase in demand under pay-as-bid pricing as bids approach marginal costs. The opposite is true for uniform pricing. Margins and, consequently, the deadweight loss, increase with demand. Thus, pay-as-bid pricing results in higher welfare for high-demand situations, whereas, potentially, uniform pricing realises higher welfare for low demand levels (in Figure 5.4, it does).

Both price rules result in the same welfare when they trade the same quantity, i.e., $x^*_{PAB} = x^*_{UP}$, which is the case for the demand level

$$\tilde{\varepsilon} = \underbrace{\frac{(\alpha - \delta_{PAB})(n + \delta_{UP})}{(\delta_{UP} - \delta_{PAB})(n + \alpha)}}_{\equiv \tilde{\mu} < 1} \bar{\varepsilon}. \quad (5.14)$$

[14] As bids are linear, the second factor increases linearly in the difference of the traded quantity to the first best quantity.

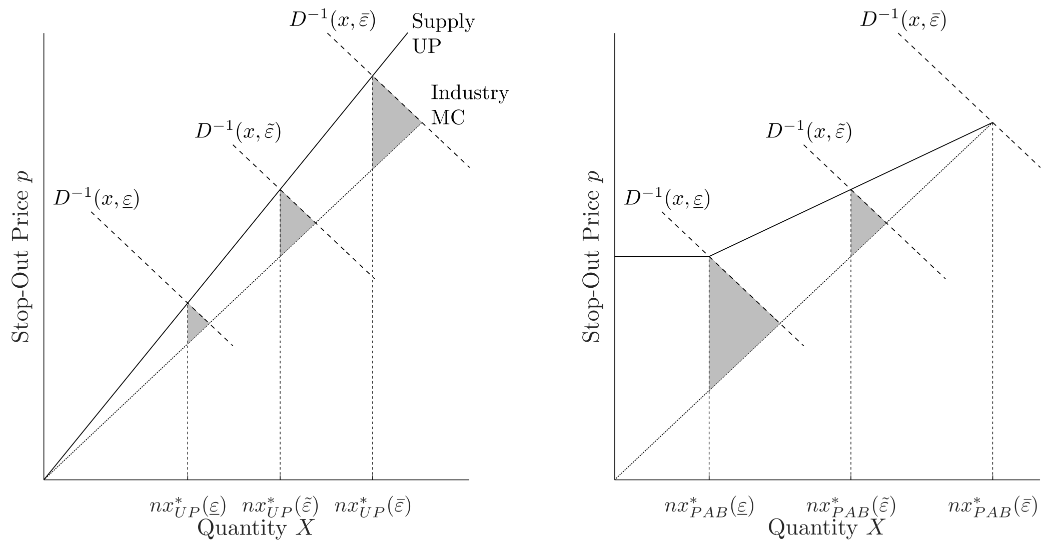

(a) Uniform pricing. **(b)** Pay-as-bid pricing.

Figure 5.4: Welfare under uniform and pay-as-bid pricing in the linear model. The dotted and solid lines represent industry-wide marginal costs and supply, respectively; the shaded triangles correspond to the deadweight loss for different demand levels.

Pay-as-bid pricing results in higher welfare if and only if $\varepsilon > \tilde{\varepsilon}$ and uniform pricing outperforms when $\varepsilon < \tilde{\varepsilon}$. Equation (5.14) shows that $\tilde{\varepsilon}$ is always strictly lower than $\bar{\varepsilon}$, which confirms that there are always some (high-demand) situations in which pay-as-bid pricing does better than uniform pricing regarding welfare. Interestingly, none of the terms on the right-hand side of equation (5.14) depend on $\underline{\varepsilon}$. Thus, we can have $\underline{\varepsilon} > \tilde{\varepsilon}$, in which case pay-as-bid pricing results in higher welfare for every single demand realisation. This is the case for any demand distribution for which the ratio of the minimum and maximum demand levels, $\mu = \underline{\varepsilon}/\bar{\varepsilon}$, is sufficiently large; i.e., that has $\mu > \tilde{\mu}$, where – according to equation (5.14) – $\tilde{\mu}$ does not depend on the demand distribution.

Expected welfare can be higher under pay-as-bid pricing even when there are some low-demand situations where uniform pricing results in higher welfare. Proposition 5.4 shows that there is a cut-off value for μ such that pay-as-bid pricing results in higher expected welfare than uniform pricing if and only if $\mu > \tilde{\mu}^E$, where $\tilde{\mu}^E > \tilde{\mu}$. An increase in μ is an increase in the lowest demand level $\underline{\varepsilon}$ for a given maximum demand level $\bar{\varepsilon}$; thus, it relatively raises the expected welfare under pay-as-bid pricing by shifting probability mass towards those cases where pay-as-bid pricing does better.

The upshot is that changes in the demand distribution affect the welfare comparison if and only if they change μ. Intuitively, μ captures demand uncertainty by measuring the distance between the minimum and maximum demand levels multiplicatively. Formally, μ measures the geometric variance of the demand level ε, or, equivalently, the variance of $\log(\varepsilon)$.[15]

Proposition 5.4 formalises these results on welfare.

Proposition 5.4 (Linear Model: Expected Welfare) *Consider the case where demand is $D(p, \varepsilon) = \varepsilon - mp$ with $\varepsilon \sim U[\underline{\varepsilon}, \bar{\varepsilon}]$, where $\underline{\varepsilon} = \mu\bar{\varepsilon}$ with $0 < \mu < 1$. There are $n \geq 2$ symmetric producers, each with the commonly known marginal cost curve $c(x) = \alpha x$ with $\alpha, m > 0$, that compete either in a pay-as-bid or a uniform-price auction. Without loss of generality, assume that $m = 1$. Furthermore, assume that the traded quantity in the pay-as-bid auction is positive, i.e., that $\mu > \underline{\mu} = \frac{(\alpha - \delta_{PAB})}{(\alpha + n)}$. If $\mu > \tilde{\mu}$, welfare is higher under pay-as-bid pricing for every single demand realisation, i.e.,*

$$\mu > \frac{(\alpha - \delta_{PAB})(n + \delta_{UP})}{(\delta_{UP} - \delta_{PAB})(n + \alpha)} \qquad \Rightarrow \qquad W^*_{PAB}(\varepsilon) > W^*_{UP}(\varepsilon) \quad \forall \varepsilon \in [\underline{\varepsilon}, \bar{\varepsilon}].$$

By contrast, a complete welfare-dominance of uniform pricing is impossible. The critical demand level for the dominance of the pay-as-bid auction, $\tilde{\mu}$, is increasing in α. For a given maximum level of demand $\bar{\varepsilon}$, the expected welfare under pay-as-bid pricing rises relative to uniform pricing when the minimum demand level increases,

[15] Thus, μ is a very specific measure of demand uncertainty; it measures the extent to which ε multiplicatively deviates from its mean. By contrast, the arithmetic variance of ε – which measures the additive deviation of ε from its mean – depends on μ *and* $\bar{\varepsilon}$. Hence, depending on $\bar{\varepsilon}$, the two pricing rules can perform differently regarding (expected) welfare for two distributions with the same arithmetic variance.

i.e., $\frac{\partial \mathbb{E}[\Delta W]}{\partial \mu} > 0$*, where* $\mathbb{E}[\Delta W] = \mathbb{E}[W^*_{PAB}] - \mathbb{E}[W^*_{UP}]$*. Thus, if the minimal demand level is sufficiently high, pay-as-bid pricing results in higher expected welfare, i.e.,*

$$\mu > \max\left\{\underline{\mu}; \tilde{\mu}^E\right\} \qquad \Rightarrow \qquad \mathbb{E}\left[W^*_{PAB}\right] > \mathbb{E}\left[W^*_{UP}\right],$$

where $\mu > \underline{\mu}$ *corresponds to the assumption of positive trade under pay-as-bid pricing and*

$$\tilde{\mu}^E = \frac{-\left(\frac{(\delta_{UP}-\alpha)^2}{(n+\delta_{UP})^2} + \frac{2(\alpha-\delta_{PAB})^2}{(n+\delta_{PAB})^2}\right) + \frac{(\delta_{UP}-\alpha)}{(n+\delta_{UP})}\sqrt{\frac{12(\alpha-\delta_{PAB})^2}{(n+\delta_{PAB})^2} - \frac{3(\delta_{UP}-\alpha)^2}{(n+\delta_{UP})^2}}}{2\left(\frac{(\delta_{UP}-\alpha)^2}{(n+\delta_{UP})^2} - \frac{(\alpha-\delta_{PAB})^2}{(n+\delta_{PAB})^2}\right)} < \tilde{\mu}. \tag{5.15}$$

If $n \leq 3$ *and* $\alpha < (4-n)/4 < m$*, we have* $\underline{\mu} > \tilde{\mu}^E$*, such that pay-as-bid pricing results in higher expected welfare for any uniform distribution with positive trade, i.e. for all* $\mu \in [\underline{\mu}, 1)$*.*

Proof. See Appendix 5.B. ■

In addition to the previously discussed, Proposition 5.4 claims that a complete welfare-dominance of pay-as-bid pricing is "more likely," i.e., can be realised for more parametrisations of the demand distribution if marginal costs are relatively flat. Similarly, it shows that pay-as-bid pricing results in higher expected welfare for *any* (uniform) distribution of demand if there are few suppliers with relatively flat marginal costs. Regarding the slope of marginal costs, Federico and Rahman (2003) obtain a similar result in a monopoly.

5.5 Implications and Discussion

When considering the policy implications of this chapter's model, it is important to acknowledge its limitations. Most importantly, the model is static. Producers do not react to market outcomes, which creates two main limitations.

First, producers' profits, which, in the model of this chapter, essentially create market inefficiency, have an important societal value in directing long-term investments (Cramton and Stoft, 2007). Willems and Yueting (2023) and Fabra *et al.* (2011) compare pricing rules regarding the resulting investment incentives with differing results.

Second, producers in the model do not condition their bidding strategies on past behaviour of their competitors; they behave as if they only interacted once. This is a critical assumption. So-called Folk Theorems imply that, by conditioning behaviour on past outcomes, less competitive behaviour than in the "one-shot" game can become an equilibrium (Mas-Collel *et al.*, 1995). Dechenaux and Kovenock (2007) and Fabra (2003) suggest that pay-as-bid pricing is less prone to such outcomes.

With these caveats in mind, the model has several important findings.

Strategic Tractability.— First, Proposition 5.1 shows that optimal behaviour under pay-as-bid pricing is tractable under fairly general conditions. Hence, fears of "tactical bidding," which were one of the reasons why the UK government's Review of Electricity Market Arrangements rejected a switch to pay-as-bid pricing (Department for Energy Security & Net Zero, 2023b), should not keep regulators from considering pay-as-bid pricing. Bidding under pay-as-bid pricing is strategically tractable and, therefore, predictable. Bringing data to the model could enable realistic welfare predictions under either pricing rule that are firmly grounded in economic theory. This way, future research could substantially lower policymakers' uncertainty regarding the effects of price rules.

Price Peaks and Consumer Surplus.— Second, uniform pricing results in higher peak prices in the model. As this holds for the general model with few functional form assumptions, it seems very likely that uniform pricing exacerbates price peaks, such as the one in January of 2025, discussed in Section 5.1. Furthermore, this chapter suggests that consumers are better off under pay-as-bid pricing.

Welfare.— However, this chapter also confirms that this increase in consumer surplus does not necessarily go along with an increase in welfare. Hence, a switch to pay-as-bid pricing is not an obvious decision. In the model, lower demand uncertainty favours pay-as-bid pricing in terms of welfare. Consequently, experiments with pay-as-bid pricing are most promising in markets with little uncertainty.

5.6 Conclusion

This chapter has derived the equilibrium strategies in the pay-as-bid auction for an oligopoly of symmetric suppliers with an arbitrary, increasing marginal cost function and an arbitrary, random, downward-sloping demand. This general model showed that peak prices are lower under pay-as-bid pricing and that both pricing rules implement the first-best welfare with zero profits for suppliers when competition is intense, because marginal costs are constant or there are infinitely many suppliers.

To derive more detailed results, this chapter has also considered a linear model, where marginal costs and demand are linear, and the intercept of the demand curve follows a uniform distribution. In the linear model, pay-as-bid pricing results in a higher expected consumer surplus because it lowers prices on inframarginal units. By contrast, the comparison of expected welfare remains ambiguous. Pay-as-bid pricing outperforms uniform pricing for at least some demand realisations. It results in higher expected welfare if demand variability is sufficiently low.

This chapter has shown that optimal bidding is tractable – and, thus, in principle, predictable – under pay-as-bid pricing under general conditions; the risk of strategic bidding should not keep regulators from considering it. A switch to pay-as-bid pricing presumably would lower peak prices and benefit consumers, whereas its effects on welfare are unclear. Positive welfare effects are most likely in markets with low demand uncertainty.

5.A Appendix: Proof of Proposition 5.3

The equilibrium under uniform pricing maximises consumer surplus conditional on the pricing rule and bids: A unit is traded if and only if the value it creates for consumers is larger than the associated price. In contrast to uniform pricing, trading a unit does not impact the price of other units such that the equilibrium quantity maximises consumer surplus. Therefore, to prove Proposition 5.3, it is sufficient to show that $CS_{PAB}(nx^*_{UP}) > CS^*_{UP}$ given that $CS^*_{PAB} > CS_{PAB}(nx^*_{UP})$. As both pricing rules would generate the same gross value to consumers when trading nx^*_{UP}, we can prove Proposition 5.3 by showing that consumers' total payments would be lower under pay-as-bid pricing when trading nx^*_{UP}, i.e., that

$$\Delta P = \underbrace{\mathbb{E}\Big[nx^*_{UP}\ \beta^*_{UP}\Big(x^*_{UP}(\varepsilon)\Big)\Big]}_{\text{Expected Payment in UP}} - \underbrace{\mathbb{E}\Big[n\int_0^{x^*_{UP}(\varepsilon)} \beta^*_{PAB}(x)\ dx\Big]}_{\text{Exp. Pay. in PAB when trading } nx^*_{UP}} > 0. \tag{5.A.1}$$

Because we do not know, for what demand levels $x^*_{UP}(\varepsilon) > x^*_{PAB}(\underline{\varepsilon})$ such that the increasing part of the bidding function $\beta^*_{PAB}(\cdot)$ becomes relevant. We can overestimate payments under pay-as-bid pricing by assuming that $x^*_{UP}(\underline{\varepsilon}) > x^*_{PAB}(\underline{\varepsilon})$, i.e.,

$$\mathbb{E}\Big[n\int_0^{x^*_{UP}(\varepsilon)} \beta^*_{PAB}(x)\ dx\Big] < \mathbb{E}\left[n\ \left(\int_0^{x^*_{UP}(\varepsilon)}\Big(\lambda + \delta_{PAB}\,x\Big)dx + 0.5\,\delta_{PAB}\ \Big(x^*_{PAB}(\underline{\varepsilon})\Big)^2\right)\right],$$

Using this bound in (5.A.1) and dividing by n, we get that for $\Delta P > 0$ it is sufficient that

$$\frac{(2\delta_{UP} - \delta_{PAB})}{2(n+\delta_{UP})^2}\,\mathbb{E}[\varepsilon^2] - \frac{2\lambda}{2(n+\delta_{UP})}\,\mathbb{E}[\varepsilon] - \frac{\delta_{PAB}\,(\underline{\varepsilon}-\lambda)^2}{2(n+\delta_{PAB})^2} > 0. \tag{5.A.2}$$

Rewriting the expected values by setting $\underline{\varepsilon} = \mu\bar{\varepsilon}$, where $\mu \in [\underline{\mu}, 1)$ with $\underline{\mu} = \lambda/\bar{\varepsilon} = \frac{(\alpha-\delta_{PAB})}{(\alpha+n)}$ by the assumption of positive trade, i.e., $\underline{\varepsilon} > \lambda$ and dividing by $(\bar{\varepsilon})^2$, we get

$$\begin{aligned}\Delta\tilde{P} = \frac{(2\delta_{UP}-\delta_{PAB})\,(1+\mu+\mu^2)}{6(n+\delta_{UP})^2} - \frac{(\alpha-\delta_{PAB})\,(1+\mu)}{2(n+\delta_{UP})\,(n+\alpha)}\\ - \frac{\delta_{PAB}}{2(n+\delta_{PAB})^2}\left(\mu - \frac{(\alpha-\delta_{PAB})}{(\alpha+n)}\right)^2 > 0,\end{aligned} \tag{5.A.3}$$

which is equivalent to (5.A.2) under the assumptions of Proposition 5.3. Given that (5.A.3) is a polynomial of degree two in μ, (5.A.3) holds if (a) $\Delta\tilde{P}(\underline{\mu}) > 0$, (b) $\Delta\tilde{P}'(\underline{\mu}) > 0$ and (c) $\Delta\tilde{P}(1) > 0$.

Proof of (a): For $\Delta\tilde{P}(\underline{\mu})$, we get

$$\frac{(2\delta_{UP}-\delta_{PAB})\,(1+\underline{\mu}+\underline{\mu}^2)}{6(n+\delta_{UP})^2} - \frac{(\alpha-\delta_{PAB})\,(1+\underline{\mu})}{2(n+\delta_{UP})\,(n+\alpha)} > 0$$

$$\overset{\substack{\delta_{PAB}<\alpha/2\\ \delta_{PAB}>(\alpha-1)/2}}{\Leftarrow} \quad \frac{(2\delta_{UP}-\delta_{PAB})}{6(n+\delta_{UP})^2} - \frac{(\alpha-\delta_{PAB})}{2(n+\delta_{UP})\,(\alpha+n)} + \frac{(2\delta_{UP}-\delta_{PAB})}{6(n+\delta_{UP})^2}\,\frac{\alpha^2}{2(\alpha+n)}\,\frac{1}{(3\alpha+2n+1)} > 0$$

$$\overset{\substack{\delta_{PAB}<\alpha/2\\ \delta_{PAB}>(\alpha-1)/2}}{\Leftarrow} \quad \frac{(2\delta_{UP}-\delta_{PAB})}{6(n+\delta_{UP})^2} - \frac{(\alpha-\delta_{PAB})}{2(n+\delta_{UP})\,(\alpha+n)} + \frac{(2\delta_{UP}-\delta_{PAB})}{6(n+\delta_{UP})^2}\,\frac{\alpha^2}{2(\alpha+n)}\,\frac{1}{(3\alpha+2n+1)} > 0$$

$$\overset{\substack{n>1\\ \delta_{PAB}<\alpha/2}}{\Leftarrow} \quad (2\delta_{UP}-\alpha/2)\,\frac{6(\alpha+n)^2+\alpha^2}{6(\alpha+n)} - 3(n+\delta_{UP})(\alpha-\delta_{PAB}) > 0$$

$$\overset{\substack{\delta_{UP}>\alpha\\ \delta_{PAB}\delta_{UP}\geq\frac{\alpha^2}{2}}}{\Leftarrow} \quad (2n-\alpha)\delta_{UP} + 3n\delta_{PAB} + a^2 - \frac{7n}{2}\alpha + \frac{\alpha^3}{4(\alpha+n)} > 0, \qquad (5.A.4)$$

To prove that (5.A.4) holds, we transform it into easy-to verify polynomials: For $n = 2$, (5.A.4) resolves to

$$\alpha^2 - 7\alpha + \frac{a^3}{2(a+2)} + (4-\alpha)\sqrt{\alpha(4+\alpha)} + 3\sqrt{9+\alpha(2+\alpha)} - 9 > 0.$$

For $n \geq 3$, use $\delta_{PAB} > \frac{(n-1)}{(2n-1)}\alpha$, to obtain the sufficient condition

$$(2n-\alpha)\delta_{UP} + \alpha^2 + \frac{\alpha^3}{4(a+n)} - \frac{(8n-1)n\alpha}{2(2n-1)} > 0. \qquad (5.A.5)$$

To reduce this problem to a simple polynomial, we can bound δ_{UP} distinguishing $2n \lesseqqgtr \alpha$. *If* $2n < \alpha$, δ_{UP} enters negatively, and we get a sufficient condition by using $\delta_{UP} < \alpha + 1$. Hence, a sufficient condition for (5.A.5) in this case, is

$$\alpha^3 + \frac{(4-14n)}{(2n-1)}\,\alpha^2 + \frac{2n(n-2)}{(2n-1)}\,\alpha + 8n^2 > 0.$$

If $2n \geq \alpha$, we can use the additional restriction on α to give a particularly tight lower bound on δ_{UP}, i.e., $\delta_{UP} > \alpha + \frac{277(n+2)}{400n(n+\alpha)}\alpha$ to rewrite the sufficient condition from (5.A.5) as

$$100n(2n-1)\alpha^2 + (-1154n^2 - 831n + 554)\alpha + 508n^3 + 1662n^2 - 1108n.$$

Proof of (b): $\Delta\tilde{P}'(\underline{\mu}) > 0$ follows directly from (a). We have,

$$\Delta\tilde{P}'(\underline{\mu}) = \frac{(2\delta_{UPA} - \delta_{PABA})(1 + 2\underline{\mu})}{6(n + \delta_{UPA})^2} - \frac{(\alpha - \delta_{PABA})}{2(n + \delta_{UPA})(n + \alpha)},$$

which is larger than $\Delta\tilde{P}(\underline{\mu})$ given that $\underline{\mu} > 0$ and $\underline{\mu}^2 < \underline{\mu}$ as $\underline{\mu} < 1$.

Proof of (c): We need to show that

$$\Delta\tilde{P}(1) = \frac{(2\delta_{UP} - \delta_{PAB})}{2(n + \delta_{UP})^2} - \frac{(\alpha - \delta_{PAB})}{(n + \delta_{UP})(n + \alpha)} - \frac{\delta_{PAB}}{2(n + a)^2} > 0.$$

$\Delta\tilde{P}(1)$ is decreasing in δ_{PAB}. Thus, $\Delta\tilde{P}(1)$ is strictly positive as long as $\delta_{PAB} < \frac{2n(\alpha+n)}{(\delta_{UP}-\alpha)}$, which is satisfied as $(\delta_{UP} - \alpha)\,\delta_{PAB} < (\alpha + 1 - \alpha)\,\delta_{PAB} = \delta_{PAB} < \alpha < 2n(\alpha + n)$.

5.B Appendix: Proof of Proposition 5.4

Critical Value for Dominance of Pay-as-Bid Pricing $\tilde{\mu}$ decreases in α.— Proposition 5.4 claims that $\frac{\partial\tilde{\mu}}{\partial\alpha} > 0$, or equivalently,

$$-(\delta_{UP} - \delta_{PAB}) + \frac{(n+\alpha)}{(n+\delta_{UP})}(\alpha - \delta_{PAB})\frac{\partial\delta_{UP}}{\partial\alpha} + \frac{(n+\alpha)}{(n+\delta_{PAB})}(\delta_{UP} - \alpha)\frac{\partial\delta_{PAB}}{\partial\alpha} < 0. \quad (5.B.1)$$

A sufficient condition for (5.B.1) to hold follows from replacing the first factors in the second and third summand by one, which, after rewriting the derivatives, gives

$$-(\delta_{UP} - \delta_{PAB}) + (\alpha - \delta_{PAB})\frac{(\delta_{UP} + n - 1)}{(2\delta_{UP} - a + n - 2)} + (\delta_{UP} - \alpha)\frac{(\delta_{PAB} + n - 1)}{(4\delta_{PAB} - a + 2n - 1)} < 0. \quad (5.B.2)$$

For $\alpha < 1/2$, use that $\frac{(n-1)}{(2n-1)}\alpha < \delta_{PAB} < \frac{\alpha}{2} - \frac{\alpha}{2(2n-1+\alpha)}$ in the first two summands, and $a/4 < \delta_{PAB} < a/2$ in the third summand. After multiplying by $2\sqrt{\alpha^2 + 2\alpha n + n^2 - 4n + 4}$, adding $(1/4 - a/2)(a^2/(2n-1)) > 0$, and multiplying by $2(2n-1)(2n-1+\alpha)$, we get that a sufficient condition for (5.B.2) is

$$\begin{aligned}\underbrace{\left(2n + \frac{1}{2}\right)}_{>0}\alpha^3 + \underbrace{\left(3n^2 - 5n + \frac{7}{2}\right)}_{>0}\alpha^2 + \underbrace{\left(-4n^3 + n^2 + 4n - 4\right)}_{<0}\alpha - 2n(n-2)^2(2n-1)\\ < \left(a^3 + (3n-5)\alpha^2 - (n-2)\alpha - n(4n^2 - 10n + 4)\right)\underbrace{\sqrt{\alpha^2 + 2\alpha n + n^2 - 4n + 4}}_{>0}.\end{aligned} \tag{5.B.3}$$

The left-hand side is negative; the first two summands are dominated by the third one given that $\alpha < 1/2$. The right-hand side is negative if and only if $n \geq 3$. Thus, (5.B.3) holds for $n = 2$. For $n \geq 3$, squaring both sides and dividing by $\alpha^2 > 0$, (5.B.3) is equivalent to

$$\begin{aligned}0 > \alpha^6 + (8n-10)\alpha^5 + \underbrace{\frac{(72n^2 - 232n + 131)}{4}}_{>0}\alpha^4 + \frac{(8n^3 - 134n^2 + 254n - 127)}{2}\alpha^3\\ + \frac{(-96n^4 + 360n^3 + 140n^2 - 804n + 447)}{4}\alpha^2\\ + (-16n^5 + 140n^4 - 264n^3 + 149n^2 + 52n - 52)\alpha\\ - (8n^5 - 176n^4 + 490n^3 - 408n^2 + 104n),\end{aligned}$$

where we can obtain an upper bound for the right-hand side by using the fact $\alpha < 1$, i.e., it is sufficient that

$$\begin{aligned}0 > \alpha^2 + (8n-10)\alpha^2 + \frac{(72n^2 - 232n + 131)}{4}\alpha^2 + \frac{(8n^3 - 134n^2 + 254n - 127)}{2}\alpha^3\\ + \frac{(-96n^4 + 360n^3 + 140n^2 - 804n + 447)}{4}\alpha^2\\ + (-16n^5 + 140n^4 - 264n^3 + 149n^2 + 52n - 52)\alpha\\ - (8n^5 - 176n^4 + 490n^3 - 408n^2 + 104n)\alpha.\end{aligned}$$

Replacing α by $2\alpha^2$ in the last two summands, after multiplying by two and dividing by α^2, it is sufficient to show that, for $\alpha < 1/2$ and $n \geq 3$,

$$0 > (8n^3 - 134n^2 + 254n - 127)\,\alpha + (-96n^5 - 192n^4 + 1084n^3 - 930n^2 + 122n + 63).$$

For $\alpha \geq 1/2$**,** a sufficient condition for (5.B.2) is $-(\delta_{UP}-\delta_{PAB})+(\alpha-\delta_{PAB})\frac{\partial \delta_{UP}}{\partial \alpha}+\frac{(\delta_{UP}-\alpha)}{2} < 0$, given that $\delta_{PAB} > (a-1)/2$, or equivalently,

$$(\alpha - \delta_{PAB}) - \left(\frac{(n-2)+2(\delta_{UP}-\delta_{PAB})+\alpha}{2}\right)(\delta_{UP}-\alpha) < 0.$$

Using the bounds $\frac{(n-1)}{(2n-1)}\alpha < \delta_{PAB} < \alpha/2$ for the first and second δ_{PAB}, respectively, we get the alternative sufficient condition

$$(n-2)^2 + \left(\frac{2n^2-n-2}{(2n-1)}\right)\alpha > (n-2)\sqrt{\alpha^2+2\alpha n+n^2-4n+4}. \tag{5.B.4}$$

For $n = 2$, (5.B.4) evidently holds, as only the positive second summand on the left-hand side is non-zero. For $n \geq 3$ both the left- and right-hand side of (5.B.4) are positive. Thus, after squaring both sides, (5.B.4) is equivalent to

$$\left(\left(\frac{2n^2-n-2}{(2n-1)(n-2)}\right)^2 - 1\right)\alpha - \frac{4}{(2n-1)} > 0.$$

Relative Expected Welfare under Pay-as-Bid Pricing Increases with Higher Minimal Demand.— The difference in welfare between the pricing rules for a given demand realisation $\varepsilon = z\bar{\varepsilon}$ with $z \in [\underline{\mu}; 1)$, where from the assumption of positive trade $\underline{\mu} = \frac{(\alpha-\delta_{PAB})}{(\alpha+n)}$, is the difference in the respective deadweight loss, i.e.,

$$\begin{aligned}\Delta W = (z\bar{\varepsilon})^2\frac{n}{2}&\underbrace{\left(\frac{1}{(n+m\alpha)}-\frac{1}{(n+\delta_{UP})}\right)}_{>0}\underbrace{\left(\frac{(\delta_{UP}-\alpha)}{(n+\delta_{UP})}\right)}_{>0}\\ &-(\bar{\varepsilon}^2)\frac{n}{2}\underbrace{\left(\frac{z}{(n+m\alpha)}-\frac{(z-\tilde{\lambda})}{(n+\delta_{PAB})}\right)}_{>0}\underbrace{\left(\frac{(\delta_{PAB}-\alpha)(z-\tilde{\lambda})}{(n+\delta_{PAB})}+\tilde{\lambda}\right)}_{>0},\end{aligned}$$

where $\tilde{\lambda} = \lambda/\bar{\varepsilon} = \frac{(\alpha-\delta_{PAB})}{(\alpha+n)}$. The derivative of ΔW with respect to μ is

$$\begin{aligned}\frac{\partial \Delta W}{\partial z} = z(\bar{\varepsilon})^2\, n&\underbrace{\left(\frac{1}{(n+m\alpha)}-\frac{1}{(n+\delta_{UP})}\right)}_{>0}\underbrace{\left(\frac{(\delta_{UP}-\alpha)}{(n+\delta_{UP})}\right)}_{>0}\\ &-(\bar{\varepsilon}^2)\frac{n}{2}\underbrace{\left(\frac{z}{(n+m\alpha)}-\frac{(z-\tilde{\lambda})}{(n+\delta_{PAB})}\right)}_{>0}\underbrace{\left(\frac{(\delta_{PAB}-\alpha)}{(n+\delta_{PAB})}\right)}_{<0}\\ &-(\bar{\varepsilon}^2)\,\frac{n}{2}\underbrace{\left(\frac{1}{(n+m\alpha)}-\frac{1}{(n+\delta_{PAB})}\right)}_{<0}\underbrace{\left(\frac{(\delta_{PAB}-\alpha)(z-\tilde{\lambda})}{(n+\delta_{PAB})}+\tilde{\lambda}\right)}_{>0} > 0.\end{aligned}$$

From $\frac{\partial \Delta W}{\partial z} > 0$, it follows that $\frac{\partial \mathbb{E}[\Delta W]}{\partial \mu} > 0$: Expected welfare can be written as

$$\mathbb{E}[\Delta W] = \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} \Delta W(\varepsilon; \bar{\varepsilon}) \frac{1}{\bar{\varepsilon} - \underline{\varepsilon}} \, d\varepsilon = \bar{\varepsilon}^2 \int_{\mu}^{1} \Delta \tilde{W}(z) \frac{1}{(1-\mu)} \, dz, \tag{5.B.5}$$

where $\Delta \tilde{W}(z) = \Delta W / (\bar{\varepsilon})^2$. Hence, $\frac{\mathbb{E}[\Delta W]}{\partial \mu} = \frac{\bar{\varepsilon}^2}{(1-\mu)} \left(\frac{1}{(1-\mu)} \int_{\mu}^{1} \Delta \tilde{W}(z) \, dz - \Delta \tilde{W}(\mu) \right) > 0$.

Conditions for Higher Expected Welfare under Pay-as-Bid Pricing.— Based on (5.B.5), pay-as-bid pricing results in higher expected welfare if and only if $\int_{\mu}^{1} \Delta \tilde{W}(z) \, dz > 0$. Given that $\frac{\partial \mathbb{E}[\Delta W]}{\partial \mu} > 0$ the critical value of μ for (5.B.6) to hold is $\tilde{\mu}^E < \tilde{\mu}$. Formally $\tilde{\mu}^E$ solves

$$\int_{\tilde{\mu}^E}^{1} \Delta \tilde{W}(z) \, dz = 0, \tag{5.B.6}$$

or equivalently, $\frac{(\delta_{UP}-\alpha)^2}{(n+\delta_{UP})^2} (1 - (\tilde{\mu}^E)^3) - \frac{(\alpha-\delta_{PAB})^2}{(n+\delta_{PAB})^2} (1 - \tilde{\mu}^E)^3 = 0$. Given that $\tilde{\mu}^E = 1$ is an obvious solution to (5.B.6) but not the desired solution with $\tilde{\mu}^E < \tilde{\mu}$, after polynomial long division with $(1 - \tilde{\mu}^E)$, we have that

$$\begin{aligned} &\left(\frac{(\delta_{UP} - \alpha)^2}{(n + \delta_{UP})^2} - \frac{(\alpha - \delta_{PAB})^2}{(n + \delta_{PAB})^2} \right) \left((\tilde{\mu}^E)^2 + 1 \right) \\ &\qquad + \left(\frac{(\delta_{UP} - \alpha)^2}{(n + \delta_{UP})^2} + \frac{2(\alpha - \delta_{PAB})^2}{(n + \delta_{PAB})^2} \right) (\tilde{\mu}^E) = 0. \end{aligned} \tag{5.B.7}$$

Distinguish two cases: If $\frac{(\delta_{UP}-\alpha)^2}{(n+\delta_{UP})^2} \geq \frac{(\alpha-\delta_{PAB})^2}{(n+\delta_{PAB})^2}$ the left-hand side of (5.B.7) is positive for all $\tilde{\mu}^E > 0$, i.e., expected welfare under pay-as-bid pricing exceeds expected welfare under uniform pricing for any μ.[16] If $\frac{(\delta_{UP}-\alpha)^2}{(n+\delta_{UP})^2} < \frac{(\alpha-\delta_{PAB})^2}{(n+\delta_{PAB})^2}$ (5.B.7) has one solution for which $\tilde{\mu}^E < 1$, which is given in (5.15).

Cases in which Expected Welfare is Higher under Pay-as-Bid Pricing for Any Level of Demand Uncertainty.— If $\frac{(\delta_{UP}-\alpha)^2}{(n+\delta_{UP})^2} > \frac{(\alpha-\delta_{PAB})^2}{(n+\delta_{PAB})^2}$ pay-as-bid pricing results in higher welfare for any $\mu \in [\underline{\mu}; 1)$. Given that $\delta_{PAB} > \frac{(n-1)\alpha}{(2n-1)}$ for this to be the case, it is sufficient that

$$\frac{(\delta_{UP} - \alpha)}{(n + \delta_{UP})} > \frac{\left(\alpha - \frac{(n-1)\alpha}{(2n-1)}\right)}{\left(n + \frac{(n-1)\alpha}{(2n-1)}\right)} = \frac{\alpha n}{(2n^2 + (\alpha - 1)n - \alpha n)} > \frac{\alpha}{(2n-1)},$$

[16] Expected welfare is higher if under pay-as-bid pricing if and only if $(1-\tilde{\mu}^E)$ times the left-hand side is of (5.B.7) is positive with $\mu = \tilde{\mu}^E$.

if $\alpha < 1$. The sufficient condition $\frac{(\delta_{UP}-\alpha)}{(n+\delta_{UP})} > \frac{\alpha}{(2n-1)}$ is equivalent to

$$\delta_{UP} > \frac{(3n-1)}{(2n-1-\alpha)}\,\alpha. \tag{5.B.8}$$

It is easy to verify that (5.B.8) holds for $\alpha < 1/2$ if $n = 2$ and $\alpha < 1/4$ for $n = 3$.

6
Conclusion

This dissertation is a contribution to microeconomic theory, particularly auction theory in procurement settings. Hopefully, it also holds some insights for real-world procurement. Chapter 2 opens the dissertation with a theoretical result: It derives an optimal procurement mechanism for a buyer trading off the procured quantity and overall spending while also targeting specific goals for both. Chapter 3 shows that this result is not as theoretical as the complexity of the resulting mechanism makes it seem. Even for such a complex mechanism, it is conceptually straightforward to design an auction that implements the optimal mechanism: Suppliers bid unit costs, and the lowest bidder supplies a quantity that depends on their bid and gets paid their ask price. Alas, these auctions are only easy to implement conceptually, not practically; determining the optimal demand schedule is not necessarily easy and requires information that real-world buyers might not have. However, Chapter 3 also proposes a promising alternative to widely used fixed-quantity auctions. Buyers can also auction off fixed budgets and procure as much as possible with that budget. Roughly speaking, Chapter 3 shows that a buyer who does not primarily care about reaching a specific quantity is better off under a budget auction than under the more traditional fixed-quantity auction. Although these budget auctions are rare, there are examples of their use. Chapter 4 can be considered a model of such a case: the Austrian network coverage procurement auction in 2020. Was it a good idea to fix the procurement budget rather than the number of communities that should get coverage? As far as it is possible to answer this question, the results of Chapter 4 suggest that it was. The chapter also serves as a reminder of how seemingly "strange" model assumptions sometimes reflect real-world situations in surprising ways. Last, Chapter 5 demonstrates how advances in economic theory can help better understand important real-world questions. Tentatively, the results of Chapter 5 suggest that one way to lower electricity prices is a switch to pay-as-bid pricing. More importantly, Chapter 5 contains much of what is needed from theory to get to a less tentative answer.

Appendices

A
Published Version of Chapter 4

The RAND Journal of Economics

ORIGINAL ARTICLE OPEN ACCESS

Auctioning Off Budgets in Procurement

Claudio Rottner

Faculty of Law, Business & Economics, University of Bayreuth, Bayreuth95447, Germany

Correspondence: Claudio Rottner (Claudio.Rottner@uni-bayreuth.de)

Received: 24 October 2025 | **Accepted:** 24 October 2025

Keywords: reverse auction | multi-unit auction | procurement | pay-as-bid

ABSTRACT

This article investigates a multi-unit pay-as-bid procurement auction where the auctioneer fixes total spending and maximizes the quantity procured with their predetermined, secret budget. Previous literature has analyzed fixed-quantity auctions, where the traded quantity is fixed but unknown to the bidders. Compared to such auctions, budget auctions lower the auctioneer's costs by introducing an additional interaction between a bidder's bids; bidders not only weigh a higher profit margin on a unit against a lower probability of supplying that unit; a higher margin on some unit also reduces the probability that the budget suffices to procure more units from the bidder.

JEL Classification: C72, D44, H57

1 | Introduction

In 2020, the Austrian telecommunications regulatory authorities (TKK and RTR) used a reverse auction to procure mobile network coverage for remote areas that had previously lacked an adequate network infrastructure (RTR 2020). In the pay-as-bid auction, each mobile network operator submitted a set of bids, each of which specified the total subsidy they demanded for providing coverage to a certain number of municipalities (TKK 2019). The auctioneer selected the combination of the different bidders' bids that maximized the number of municipalities to which coverage would be provided, subject to the constraint that the total subsidy did not exceed a secret, predetermined budget (TKK 2019).

The Austrian spectrum auction can be considered the real-world equivalent of the model in this article. The model describes a novel strategic effect in bidding behavior. Due to the auctioneer's budget cap, there is an interaction between a bidder's bids on different units: A high bid on some unit reduces the residual budget. Thus, fewer of the bidder's bids can be considered. Such a link between bids on different units does not have a correspondence in previous literature.

This article derives the symmetric equilibrium in a budget auction when bidders have the same cost function and share a stochastic belief about the unknown budget of the auctioneer. Previous literature (e.g., Holmberg 2009; Pycia and Woodward 2023) has considered the case where the auctioneer fixes the traded quantity, which is unknown to the bidders. In a pay-as-bid procurement auction in such a setting, when bidders raise their bids, they trade off the additional margin on a particular unit (price effect) with a decreased probability of realizing the margin on this unit (quantity effect). By contrast, increasing the profit margin on some unit also hurts the chances of realizing the profit margin on *other* units in the budget auction. Higher bids decrease the residual budget and therefore the probability that bids on higher quantities can be served without exceeding the auctioneer's budget (budget effect). Consequently, the auctioneer's costs in a budget auction are strictly lower than in a comparable fixed-quantity auction, that is, a fixed-quantity auction with the same distribution of realized demand in terms of quantities. The cost advantage of the budget auction is most pronounced when traditional fixed-quantity auctions suffer from soft competition, that is, when there are few bidders whose costs exhibit strong diseconomies of scale.

This article proceeds as follows. Section 2 discusses related literature. Section 3.1 first explains the basic model setup for the fixed-quantity and the budget auction. For ease of exposition, it then briefly discusses the symmetric equilibrium in the fixed-quantity auction translating the model from Pycia and Woodward (2023) to the procurement context. Afterward, I characterize the symmetric equilibrium in the strategically more complex budget auction. Section 4 shows that the budget auction comes at a lower cost to the auctioneer than a comparable fixed-quantity auction. Section 5 discusses some of this article's limitations before Section 6 concludes.

2 | Related Literature

Traditionally, auction theory assumes that the traded quantity is fixed in advance. In the "single-unit", or rather, "'single-winner" context, the analysis of auctions with an endogenous traded quantity dates back at least to Hansen (1988). Dasgupta and Spulber (1989) demonstrate that variable-quantity auctions are optimal under plausible assumptions, both in cases with a single winner and multiple winners. However, in contrast to typical real-world auctions or budget auctions, implementing these optimal auctions requires the auctioneer to possess extensive knowledge regarding the bidders' cost structure. Budget auctions, as a specific form of auctions with a variable traded quantity, have been discussed in the single-winner context by Dastidar (2008), Deck and Wilson (2008), and Liu and Parlour (2014).

However, this article is better understood against the backdrop of the literature explicitly characterizing equilibria in multi-unit (multi-winner) auctions. Burkett and Woodward (2020b) are a notable outlier in this literature because they consider a *discrete* uniform-price auction with *uncertainty over individual bidder types*. The rest of this literature, including this article, rather follows Klemperer and Meyer (1989) in analyzing multi-winner auctions with a continuous traded quantity and bidders with an identical, known cost structure. Specifically, Klemperer and Meyer (1989) consider a uniform-price procurement auction in which the bidders share a belief about the random, downward-sloping demand function. They explicitly characterize the potentially infinite number of symmetric equilibria. They show that the equilibrium bidding functions lie in a band that bounds them from above and below. If one were to adapt their model such that demand is completely inelastic, that is, the traded quantity is fixed in advance, this would correspond to the uniform-price equivalent of the model of a pay-as-bid fixed-quantity auction in this article. In such a model, the infinitely many symmetric equilibria can be arbitrarily unattractive to the auctioneer. These equilibria are called the "'low-revenue equilibria" of the uniform price auction. Their occurrence was first observed by Wilson (1979).

For the pay-as-bid auction, Holmberg (2009) derives the symmetric equilibrium for the reverse auction with completely inelastic demand. Just like Klemperer and Meyer (1989), he assumes that bidders have an identical, known cost function and share a belief over the distribution of demand. However, for his model to work, he requires that demand surpasses the suppliers' production capacity with a positive probability, which limits the applicability of his results. Pycia and Woodward (2023) relax this assumption in the context of a forward auction. Their model forms the fixed-quantity baseline to which I compare the budget auction in this article.

Pycia and Woodward (2023) also consider an endogenous traded quantity in their analysis of random reserve prices. Random reserve prices punish high bids by increasing the probability that the increased bid will not be considered. This reinforces the quantity effect and is, therefore, possible in both fixed-quantity and budget auctions. Thus, the insights into secret reserve prices complement the analysis in this article, which focuses on a novel strategic effect deterring high bids by reducing the probability that bids on other units are accepted.

For the uniform-price auction, some more articles consider a variable traded quantity besides Klemperer and Meyer (1989). These models address the issue of low-revenue equilibria by allowing the auctioneer to restrict the traded quantity after observing the bids (e.g., Lengwiler 1999; Back and Zender 2001; McAdams 2007).[1] Like in this article, the threat to reduce the traded quantity if bids are too unattractive to the auctioneer prompts the bidders to reduce bid shading. However, the auction designs from this literature can only be applied in the real world under exceptional circumstances. Consider McAdams (2007), for example. In his model, the auctioneer adjusts supply *at will* ex-post, that is, does not make *any* commitment regarding the traded quantity, and the bidders reveal their type *truthfully* as a consequence.

The auctioneer could exploit this design by learning the bidders' types without trading any significant quantity and then extracting rents from the bidders in trade outside the auction. According to Rothkopf et al. (1990), the bidders' unwillingness to reveal their type is one of the main reasons for the scarcity of second-price auctions in the real world. This concern is mitigated in the pay-as-bid setting considered here. Furthermore, the auctioneer in this article does commit to a budget, although the specific amount is not disclosed. Bergemann and Hörner (2018) and Allen et al. (2024) suggest that such opacity regarding parts of the auction design is not unusual in real-world auctions.

This raises the question of how the auctioneer can credibly commit when the bidders cannot observe this commitment. Tillio et al. (2016) argue that credible commitment is possible if the auctioneer is indifferent between the available alternatives. However, this seems unlikely in most applications and is certainly not the case for the budget auction. This suggests other mechanisms by which the auctioneer can make their private commitment credible exist. These include trust, as in the case of the Austrian coverage auction, and trusted intermediaries, such as eBay, for secret reserve prices on the platform. Theoretically, there are also trustless solutions. For example, the auctioneer could publish the budget in an encrypted document before and the encryption key after the auction. However, it seems unlikely that procurement agencies would opt for such a solution. Equally theoretically, uncertainty over the budget could also stem from tying the budget to a random but ex-post-verifiable outcome, for example, proceeds from future auctions of emission allowances. Why the auctioneer might *want* to implement an uncertain budget is discussed in Section 5.

3 | Model

3.1 | Setup

The auctioneer wants to procure multiple units of a homogeneous good using a pay-as-bid auction. The main claim of this article is that the auctioneer is better off when doing so using a budget auction than when using a fixed-quantity design, which has been analyzed by Holmberg (2009) and Pycia and Woodward (2023).

3.1.1 | Bidders

For both auction designs, assume that all of the $n \geq 2$ bidders have the same marginal cost $c(x) > 0$ for producing the x-th unit with $0 < c'(x) < \infty$.[2] The bidders know each others' costs, whereas the auctioneer does not. The auction rules require each bidder to specify a weakly positive, twice differentiable, finite, and increasing bidding function $b_i(x)$, which denotes the marginal payment bidder i demands for providing their x-th unit. I assume that the submitted bid functions are smooth in the sense that they have *finite* first and second derivatives for $x > 0$.[3] The inverse of $b_i(x)$ is denoted by $\beta_i(\cdot)$. The auctioneer pays bidder i $B_i(x) = \int_0^{x_i} b_i(y)\,dy$ when buying x_i units from them. I refer to $B_i(x)$ as bidder i's cumulative bids. The pay-off of bidder i when assigned quantity x is

$$\pi_i(x; B_i) = B_i(x) - \int_0^x c(y)\,dy.$$

3.1.2 | Fixed-Quantity Auction

In the fixed-quantity auction, the auctioneer procures X units, whatever the costs of doing so are. The auctioneer minimizes the total payments to bidders subject to the constraint that X units are procured. In line with previous literature, assume that the bidders do not know the quantity the auctioneer wants to procure but believe that it is distributed on some interval $[0, \bar{X}]$ according to the distribution function $G(X)$ with $G'(X) = g(X) > 0$ and $\bar{X} < \infty$. This setup for the fixed-quantity auction is the reverse auction equivalent of the model from Pycia and Woodward (2023).

3.1.3 | Budget Auction

In the budget auction, the auctioneer intends to procure as many units of the good as possible with the fixed budget A. Given the submitted bids, the auctioneer solves the following optimization problem to determine the quantity x_i they buy from bidder i:

$$\begin{aligned} \max_{x_1,\ldots,x_n} \quad & \textstyle\sum_{i=1}^n x_i \\ \text{s.t.} \quad & \textstyle\sum_{i=1}^n B_i(x_i) \leq A \\ & \forall i : x_i \geq 0. \end{aligned}$$

Assume that each bidder submits bids at least up to the quantity that the auctioneer wants to assign to them. The first-order conditions of this problem imply that in an interior solution with $x_i > 0$ for all bidders, all bidding functions must have the same value at their respective chosen quantity. Refer to this unique value as the market (clearing) price. Consequently, with increasing bids, the auctioneer procures all units for which the demanded price is lower or equal to the market price.

The auctioneer's budget is unknown to the bidders. They only share a common belief that the auctioneer's budget follows a twice-differentiable cumulative distribution $F(A)$ on $[0, \bar{A}]$ with density $f(A) > 0$, $\bar{A} < \infty$ and $f(\bar{A}) < \infty$, as well as $f'(\bar{A}) < \infty$. The bidders' beliefs are rational in the sense that $A \in [0, \bar{A}]$.

3.2 | Symmetric Equilibrium in the Fixed-Quantity Auction

Pycia and Woodward (2023) have extensively analyzed the fixed-quantity auction. Proposition 1 summarizes their result on equilibrium behavior.

Proposition 1. *(**Pycia and Woodward** 2023) The symmetric equilibrium in a pay-as-bid fixed-quantity auction, in which n suppliers with costs $c(x)$ believe that the demanded quantity is distributed according to the distribution function $G(\cdot)$ on $[0, \bar{X}]$, is described by*

$$b(x) = c(x) + [1 - G(nx)]^{\frac{1-n}{n}} \int_x^{\bar{X}/n} c'(y)\,[1 - G(ny)]^{\frac{n-1}{n}}\,dy.$$

Bids are strictly above cost for all but the last unit and equal cost for the last unit.

To better understand the differences between the fixed-quantity and the budget auction regarding equilibrium behavior, the following lays out the crucial steps in the derivation of Proposition 1.

As bids are increasing, there is a unique market clearing price p, which ensures that demand equals supply, that is, $\sum_i \beta_i(p) = X$, and all units for which the bids are below p are supplied. The quantity x_i supplied by player i is implicitly defined by $p = b_i(x_i)$. Following Pycia and Woodward (2023), the probability that player i supplies weakly less than x units given some set of strategies $\boldsymbol{b} = (b_i, \boldsymbol{b_{-i}})$ is

$$\text{Prob.}(x_i \leq x; \boldsymbol{b}) = G\left(x + \sum_{j \neq i} \beta_j(b_i(x))\right).$$

Bidder i winning some quantity x implies a market clearing price of $b_i(x)$, which means that the other bidders supply $\sum_{j \neq i} \beta_j(b_i(x))$ units. Consequently, overall demand must not exceed $x + \sum_{j \neq i} \beta_j(b_i(x))$ if player i is to win x units or less. As shown by Pycia and Woodward (2023), this probability distribution allows us to write the expected profits of bidder i as

$$\max_{b_i(x)} \quad \mathbb{E}[\pi_i] = \int_0^{\bar{x}_i} [b_i(x) - c(x)] \left[1 - G\left(x + \sum_{j \neq i} \beta_j(b_i(x))\right)\right] dx, \tag{1}$$

where $\bar{x}_i = \inf\{x : G(\cdot) = 1\}$. The intuitive interpretation of Equation (1) is that the profit margin on a particular unit must be weighted by the probability with which it is realized. Due to the pay-as-bid nature of the auction, the profit margin on some quantity x is realized whenever the bidder wins x or more units.

As there is (abstracting away from the monotonicity constraint on bids) no interaction between the bids on different quantities, the first-order condition to the bidder's problem (1) is obtained by point-wise maximization of the integrand. Holmberg (2009) and Pycia and Woodward (2023) derive the first-order condition as

$$\left(\underbrace{1 - G(X)}_{\text{Price Effect}}\right) - \underbrace{[b_i(x) - c(x)]\, g(X) \sum_{j \neq i} \beta_j'(b_i(x))}_{\text{Quantity Effect}} = 0, \quad (2)$$

where the market clearing condition $x_i + \sum_{j \neq i} \beta_j(b_i(x)) = X$ has been used.[4]

The first-order condition in Equation (2) demonstrates the trade-off between price and quantity effect. The first term represents the positive effect of raising bids. In all cases where the bidder wins more than x units, they benefit from additional revenue if they increase $b_i(x)$ (price effect). The second term shows that an increase in the bid on unit x reduces the probability that the corresponding profit margin is realized (quantity effect). This effect constitutes the competition aspect of the auction. If a bidder demands a higher payment for some unit, the auctioneer considers more of the other bidders' bids first. This reduces the probability that the bidder gets to supply the unit in question.

Given that the quantity assigned to bidder i is strictly increasing in total demand X, Equation (2) allows us to gain further insights into optimal bidding behavior. As long as total demand is lower than $\bar{X}$, the first term in Equation (2) is positive, such that the first-order condition can only be satisfied if the corresponding bids are above marginal costs. By the same argument, for the highest quantity ever won, which is the quantity $\bar{x}_i$ won when $X = \bar{X}$, it must be that the corresponding bid equals marginal cost, as $1 - G(\bar{X}) = 0$. Intuitively, the price effect of increasing the bid on the last unit is zero, as the additional margin is only realized when $X = \bar{X}$. By contrast, the risk of losing the (profitable) unit when increasing the bid remains.

Thus, we obtain the terminal condition $b_i(\bar{x}_i) = c(\bar{x}_i)$. Taken together with symmetry, that is, $b_i(x) = b_j(x)\, \forall i, j$, this allows us to write the symmetric equilibrium, based on Equation (2), as

$$b(x) = (n-1)[1 - G(nx)]^{\frac{1-n}{n}} \int_x^{\bar{x}} [1 - G(ny)]^{-1/n} g(ny)\, c(y)\, dy \quad (3)$$

where $\bar{x}_i = \bar{x} = \bar{X}/n$. Integration by parts results in the formulation of Proposition 1.

3.3 | Symmetric Equilibrium in the Budget Auction

Next, consider the symmetric equilibrium in the budget auction.

3.3.1 | Bidder's Problem

Given a strategy profile $\boldsymbol{B} = (B_i, \boldsymbol{B}_{-i})$, for bidder i to win x units or less, it must be that the budget suffices at most to pay for the x units of bidder i and the $\sum_{j \neq i} \beta_j(b_i(x))$ units that the other bidders will supply in this situation, that is,

$$\text{Prob.}(x_i \leq x; \boldsymbol{B}) = F\left(B_i(x) + \sum_{j \neq i} B_j\left(\beta_j(b_i(x))\right)\right).$$

As before, the bidder's expected profit is given by weighting the profit margin on some unit x with the probability that the bidder wins at least x units. The problem of bidder i is

$$\max_{B_i(x)} \mathbb{E}[\pi_i] = \int_0^{\bar{x}_i} \underbrace{[b_i(x) - c(x)]\left[1 - F\left(B_i(x) + \sum_{j \neq i} B_j\left(\beta_j(b_i(x))\right)\right)\right]}_{\equiv \Psi(B_i, b_i, x)} dx \quad (4a)$$

$$B(0) = 0 \quad (4b)$$

$$B_i(\bar{x}_i) = \bar{A} - \underbrace{\sum_{j \neq i} B_j\left(\beta_j(b_i(\bar{x}_i))\right)}_{\equiv \phi(b_i, \bar{x}_i)}. \quad (4c)$$

The bidder's best response problem comes with two boundary conditions. First, by construction, cumulative bids start at zero. Second, the interval of relevant quantities is endogenously determined by the considered bidder's bids. As bids are increasing, the highest quantity a bidder can win is the quantity the bidder wins when the auctioneer auctions off the highest possible budget $\bar{A}$.

3.3.2 | Strategic Effects

A closer look at Equation (4a) reveals that the bid on some unit x is chosen according to three considerations:

1. Its effect on the profit margin on unit x (price effect).
2. Its effect on the probability of realizing the profit margin on unit x (quantity effect).
3. Its effect on the probability of winning units $y > x$ (budget effect).[5]

The price effect is reflected in the first bracket of the integrand in (4a). The quantity effect is described by the second summand in the distribution function. It reflects competition in the sense that it determines the order in which the different bidders' bids are considered by the auctioneer. As was seen in Section 3.2, the trade-off between price and quantity effect characterizes the equilibrium in the fixed-quantity auction. The budget effect, which is reflected in the $B_i(x)$ in the distribution function, has no correspondence in the fixed-quantity auction. To make the distinction between the effects clear, think of bids and cumulative bids as independent variables. In this case, demanding a higher total compensation does not contribute to profits. However, it reduces the residual budget and thus makes it less likely that the budget suffices to buy further units from the bidder. Interestingly, in contrast to the quantity effect, this negative effect of raising bids is not a competition effect. It is not about the relative attractiveness of the different bidders' bids. Rather, higher cumulative bids of a bidder shift the probability of winning higher quantities downward independently of what the other bidders do.

3.3.3 | Equilibrium Bidding Behavior

As the bidder's problem from (4a) includes an interaction between the bids on different units, the first-order condition for this variational calculus problem is (see, e.g., Chiang 1992)

$$\frac{\partial \Psi}{\partial B_i} - \frac{d}{dx}\frac{\partial \Psi}{\partial b_i} \stackrel{!}{=} 0. \tag{5}$$

Roughly speaking, $\partial\Psi/\partial B_i$ represents the budget effect, $\partial\Psi/\partial b_i$ is the balance of price and quantity effect, and the d/dx accounts for the relation between bids and cumulative bids. The relevant derivatives can be found in Appendix A.1. Using symmetry, that is, $b(x) = b_i(x) = b_j(x)\ \forall i, j$, the Euler equation emerges as the third-order differential equation[6]

$$b''(x) = \frac{[b'(x)]^2}{b(x)[b(x)-c(x)]}\Big[3b(x) - \frac{(n-2)}{(n-1)}c(x) - \frac{c'(x)}{b'(x)}b(x) + n\frac{f'(nB(x))}{f(nB(x))}[b(x)-c(x)]\frac{[b(x)]^2}{b'(x)}\Big]. \tag{6}$$

The two boundary conditions in (4b) and (4c) eliminate two degrees of freedom. The initial condition (4b) removes the first degree of freedom. The second piece of information lies in the terminal condition from (4c). For the symmetric case, this implies

$$B(\bar{x}) = \frac{\bar{A}}{n}. \tag{7}$$

However, Equation (7) only pins down the level of cumulative bids at the last unit but does not tell us what the last unit is. Appendix A.2 shows that the corresponding transversality condition implicitly defines the last unit through $b_i(\bar{x}_i) = c(\bar{x}_i)$. Equivalently, for the symmetric case, where $\bar{x} \equiv \bar{x}_i = \bar{x}_j$, we have

$$b(\bar{x}) = c(\bar{x}). \tag{8}$$

Equation (8) says that the last unit should be chosen in such a manner that all opportunities to realize additional profits are exhausted; that is, there is no positive margin on the last unit. Put differently, of the three strategic effects discussed earlier, only the price effect favors higher bids. However, as the weight bidders attach to the profit margin on the last unit is zero, it is absent for the last unit. The only effect active on the last unit is the quantity effect. Thus, the last unit should be offered at marginal cost.

Finally, Appendix A.2 shows that the smoothness assumption on $b(x)$ implies

$$b'(\bar{x}) = \frac{n-1}{2n-1}c'(\bar{x}). \tag{9}$$

Any increasing continuation of the bidding function beyond $\bar{x}$ stabilizes the equilibrium.[7] Proposition 2 summarizes the findings on the symmetric equilibrium.

Proposition 2. *The symmetric equilibrium in a pay-as-bid budget auction, in which n suppliers with costs $c(x)$ believe that the budget is distributed according to the distribution function $F(\cdot)$ on $[0, \bar{A}]$, is described by Equations (4b) and (6)–(9). Bids are increasing and strictly above cost for all but the last unit. Bids equal costs for the last unit.*

In addition to the previously discussed, Proposition 2 claims that bids are above marginal cost and satisfy the monotonicity constraint imposed in the model setup. The corresponding proofs and the second-order conditions are relegated to Appendices A.3 and A.4, respectively. Generally, there is no closed form for the symmetric equilibrium.

3.3.4 | Example

The following example, which allows for a closed-form representation, will later serve to illustrate the cost advantage of the budget auction. Consider the case where $n = 2$, $c(x) = \alpha x + \gamma$, and $A \sim U[0, \bar{A}]$ with $\alpha, \gamma, \bar{A} > 0$. The symmetric equilibrium is

$$b(x) = \frac{\sqrt{36\gamma^2 + 60\alpha\bar{A}} + 9\gamma}{15} + \frac{\alpha}{3}x.$$

The second-order condition for this example is discussed in Appendix A.4.

4 | Cost Advantage of the Budget Auction

Auctioning off a secret budget instead of an unknown quantity introduces an additional consideration into the bidders' decision. Under both auction designs, when raising bids, the bidders need to weigh additional profits on a unit (price effect) against a lower probability of realizing them (quantity effect). However, in the budget auction, increasing the bid on some quantity has the additional adverse effect of making it less likely to win higher quantities, as the auctioneer runs out of budget "'earlier" (budget effect). Therefore, intuitively, the budget auction should be more attractive to the auctioneer.

Testing this hypothesis requires fully specifying the distributions in which bidders believe under the two auction designs, respectively. Luckily, the bidders' knowledge of equilibrium bids allows them to seamlessly move between expectations over total spending and the procured quantity. Given their equilibrium bids in the budget auction, the bidders can map every potential budget into a quantity, which the auctioneer procures with that budget in total. Thus, the bidders can directly translate the probability distribution over budgets, which describes their beliefs in a budget auction, into a "'corresponding" probability distribution over the number of units procured by the auctioneer. This corresponding distribution describes the bidders' belief in a comparable fixed-quantity auction.

For the case of symmetric bidders, the number of units X the auctioneer procures in a budget auction with a given budget is implicitly given as $A = n\,B^B(X/n)$ (where the superscript B denotes the budget auction). Thus, the probability that the auctioneer procures less than X units is equivalent to the probability that the budget does not exceed $nB^B(X/n)$. Therefore, the bidders believe that the auctioneer's demand follows the distribution

$$\tilde{G}(X) \equiv F\big(n\,B^B(X/n)\big) \tag{10}$$

in the corresponding fixed-quantity auction. These corresponding auctions can also be generalized to bidders who do not bid the same in equilibrium because they are asymmetric. In this case, we can write the budget needed to procure a given total quantity through a budget auction as a function of the quantity x_i supplied by bidder i. We then get

$$\tilde{G}(X) \equiv F\left(B_i^B(x_i) + \sum_{j \neq i} B_j^B(\beta_j^B\left(b_i^B(x_i)\right)) \right) \tag{11}$$

with the implicit definitions $A = B_i^B(x_i) + \sum_{j\neq i} B_j^B(\beta_j^B(b_i^B(x_i)))$ and $x_i = X - \sum_{j\neq i} \beta_j^B(b_i^B(x_i))$. This gives $\tilde{g}(X) = f(\cdot)\frac{\partial A}{\partial x_i}\frac{\partial x_i}{\partial X} = f(\cdot) b_i^B(x_i)$ because a marginal change in the procured quantity must be scaled by the corresponding payment. Comparing equilibrium bids in a budget auction with the corresponding fixed-quantity auction yields two results. First, directly formalizing the intuition from the previous section, when fixing the other bidders' bids, a bidder's best response is to bid higher in in the fixed-quantity auction than in the budget auction (Proposition 3). Second, for symmetric bidders, this also entails lower equilibrium bids in the budget auction (Proposition 4).

Proposition 3. *Consider a fixed-quantity auction with demand normalized according to Equation (11). Fix the bids of bidders $j \neq i$ in the fixed-quantity auction at the level of their equilibrium bids in the corresponding budget auction, that is, $b_j(x) = b_j^B(x)$. Then, some bidder i's best response is to bid $b_i(x) > b_i^B(x)$ for $x < \bar{x}_i$.*

Proof. Consider the best response $b_i^B(x)$ of bidder i in the budget auction. It has them bidding truthfully on the last unit they receive (see Appendix A.2). With reference to Equation (A2), this means that $\frac{\partial \Psi}{\partial b_i^B}\Big|_{x_i=\bar{x}_i} = 0$. This formalizes the intuition that the optimal bid equalizes price and quantity effect because the budget effect is absent for the last unit. However, the bidder's best response has him bid below the level that would balance these two effects for all other units, that is, $\forall x < \bar{x}_i : \frac{\partial \Psi}{\partial b_i^B} > 0$. Formally, the bidder's best response entails bidding above costs for these units (see Appendix A.3). Thus, there is a budget effect in the sense that $\forall x < \bar{x}_i : \frac{\partial \Psi}{\partial B_i^B} < 0$, as can be seen from Equation (A1). Therefore, by the first-order condition (5), $\forall x < x_i : \frac{d}{dx}\frac{\partial \Psi}{\partial b_i^B} < 0$. It follows that along the equilibrium path $\forall x < \bar{x}_i : \frac{\partial \Psi}{\partial b_i^B} > 0$. The normalization of demand implies that $\frac{\partial \Psi}{\partial b_i^B}$ is equivalent to the derivative in the first-order condition of the bidder's problem in the fixed-quantity auction if evaluated at $b_i(x) = b_i^B(x)\, \forall i$. Intuitively, the normalization ensures that the density attached to every unit is the same under both auction designs such that price and quantity effects are identical. If the bidder's problem in the fixed-quantity auction is well behaved, that is, concave in b_i, the level of their best response bids is above $b_i^B(x)$. □

Although Proposition 3 demonstrates the incentive wedge between comparable budget and fixed-quantity auctions, it does not have direct implications for the equilibrium outcomes. Proposition 4 entails those at the cost of restricting attention to symmetric equilibria. It shows that equilibrium bids in a budget auction, denoted by $b^B(x)$, are lower than the equilibrium bids, $b^{FQ}(x)$, in a comparable fixed-quantity auction, where demand is normalized according to Equation (10).

Proposition 4. *If demand is normalized according to (10), bids in the symmetric equilibrium are lower in the budget auction than in the comparable fixed-quantity auction:*

$$b^B(x) \leq b^{FQ}(x) \quad \forall x \in [0, \bar{x}],$$

where the inequality is strict for $x \in [0, \bar{x})$ and the superscripts FQ and B denote the fixed-quantity and budget auction, respectively.

The result of Proposition 4 is very strong in that it postulates that *every single unit* is less expensive in the budget auction.

The equality of bids for the last unit follows from the fact that the maximum number of units an individual bidder supplies is identical between the corresponding auctions by construction. As bidders do not have a positive margin on the last unit under both auction designs, the bids on the last unit are identical in the corresponding auctions.[8] The proof of the remainder of Proposition 4 is relegated to Appendix B. A sufficient condition for higher equilibrium bids in the fixed-quantity auction is that $\frac{\partial \Psi}{\partial b_i}\Big|_{B_i=B_j=B^B, b_i=b_j=b^B} > 0$ for $x < \bar{x}$. This means that in a hypothetical situation without budget effect, bids should be higher than in the equilibrium of the budget auction. As the budget effect pulls toward lower bids, this is the case. Formally, the proof of this is analogous to the one in Proposition 3.

4.1 | Example

Let us revisit the example from Section 3.3. Equilibrium bids in the budget auction with $n = 2$, $c(x) = x$, and $A \sim U[0, 1]$ are $B^B(x) = 1/6\, x^2 + \sqrt{60}/15\, x$. The probability that less than X units are procured, that is, the corresponding distribution in the fixed-quantity auction is[9]

$$\tilde{G}(X) = \text{Prob.}(A \leq 2\, B^B(\frac{X}{2})) = \frac{1}{12}X^2 + \frac{\sqrt{60}}{15}X.$$

The left panel of Figure 1 plots the difference between the bids in the two auction designs, that is, $b^{FQ}(x) - b^B(x)$. Bids converge in the quantity supplied. Intuitively, the budget effect, which explains the difference between bids in the first place, becomes less relevant with increasing quantities. Both margins and the weight attached to them decrease in the quantity supplied. Consequently, the budget effect gives less and less reason to lower bids to increase the probability of realizing the margins on higher units as the quantity increases.

The right panel of Figure 1 shows the auctioneer's total cost of procurement in the budget auction as a share of the total procurement costs in the corresponding fixed-quantity auction, that is, $\frac{n\, B^B(X/n)}{n\, B^{FQ}(X/n)}$. In this example, the procurement costs in the budget auction are between roughly 96% and 99% of the procurement costs in the fixed-quantity auction.

Evidently, the size of the cost advantage of the budget auction varies with the model parameters. In particular, we are inter-

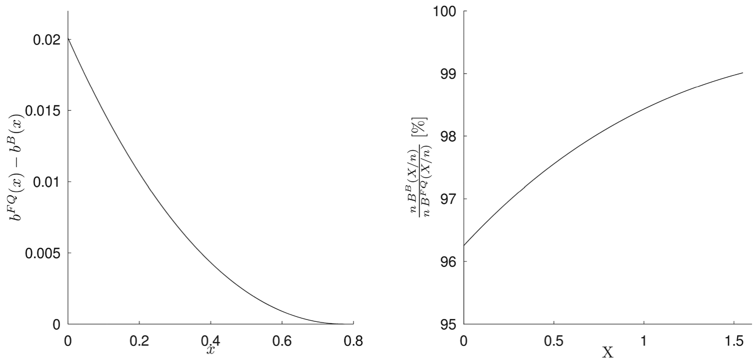

FIGURE 1 | The difference between the equilibrium bids in the fixed-quantity auction and the equilibrium bids in the budget auction to which it corresponds ($n = 2$, $c(x) = x$ and $A \sim U[0,1]$). On the left: difference in bids, on the right: total costs of procurement in the budget auction as a share of the total costs in the corresponding fixed-quantity auction.

ested in the impact that some specific model parameter, say n, has on the relation of total procurement costs under the two auction designs; that is, we would like to know $\frac{d}{dn}\frac{B^B(x)}{B^{FQ}(x)}$. This requires gauging the relative size of $\frac{d}{dn}B^B(x)$ and $\frac{d}{dn}B^{FQ}(x)$ for $\forall x < \bar{x}$. However, the established approaches to doing analytic comparative "'statics" on the equilibrium paths do not allow to do so. Even though Oniki (1973) generalizes comparative statics via implicit differentiation of the first-order condition to the dynamic context, his approach only allows *signing* the relevant expressions by drawing a phase diagram of the corresponding system of linear differential equations. Similarly, the primal-dual methodology, comprehensively set out in Caputo (2005), uses the information contained in the second-order conditions to *sign* the effect of parameter changes on the optimal paths *integrated over the entire planning horizon*, that is, $\int_0^{\bar{x}} \frac{d}{dn}B^{B/FQ}(x)\,dx$ in the case at hand. Thus, these approaches do not permit extracting the necessary information on the relative size of the impact of a parameter between the two auction designs in the absence of a closed-form solution. Consequently, the following only provides suggestive evidence on the economic contexts favorable to the budget auction based on numeric calculation.

The budget auction reduces competition problems faced by fixed-quantity auctions. Put differently, the cost advantage of the budget auction should be particularly large when fixed-quantity auctions suffer from soft competition. The general intuition is that bids in the budget auction are reduced to increase the probability of realizing margins on other units. As a consequence, bids are reduced most strongly when the margins on these other units are high, that is, competition is soft. Specifically, the budget auction should have a stronger cost advantage over the fixed-quantity auction when there are few bidders whose cost structure exhibits strong diseconomies of scale.

Figure 2 illustrates these points based on numerical calculation. The findings are robust to various specifications regarding the form of costs and the underlying distribution. To numerically construct a measure of the budget auction's cost advantage, we first need to determine equilibrium bids in the budget auction. To do this, we can guess an arbitrary value for $\bar{x}$, which fixes $B(\bar{x})$, $b(\bar{x})$, $b'(\bar{x})$ and $b''(\bar{x})$ by Equations (7), (8), (9), and (A8), respectively. Using the Euler equation from (6), we can develop the solution "'backward," that is, calculate bids for ever lower quantities. The correct choice of $\bar{x}$ is identified by forcing $B(0) = 0$. We can then numerically determine the distribution $\tilde{G}(X)$, and equilibrium bids in the corresponding fixed-quantity auction.

The left panel of Figure 2 plots the cost advantage, that is, $\frac{n\,B^B(X/n)}{n\,B^{FQ}(X/n)}$, for the previous example not only for the case of two (solid line) but also four and ten bidders. As can be seen, the fewer bidders there are, the lower the number of units the auctioneer can procure with a given budget.[10] Under both auction designs, a decreasing number of players softens competition and thus increases margins. Increased margins mean that the budget effect becomes more relevant. Therefore, the relative cost advantage of the budget auction increases when there are fewer bidders.

For the same reason, stronger diseconomies of scale increase the cost advantage of the budget auction, as can be seen in the right panel of Figure 2.[11] Again, the solid line corresponds to the previous example, and solutions are obtained numerically. To understand why diseconomies of scale increase profit margins, first consider the case of a flat marginal cost curve. In this case, the considered auctions would collapse into Bertrand competition with bids competed down to cost. It is only when marginal costs are increasing that bidders can realize profits because other bidders who already supply a larger quantity cannot exert full competitive pressure, as they are in a less favorable cost position. Extending this logic, the faster costs increase, that is, the larger the diseconomies of scale, the higher the profit margins bidders can realize. As margins increase, the budget effect becomes more relevant, and the cost advantage of budget auctions increases.

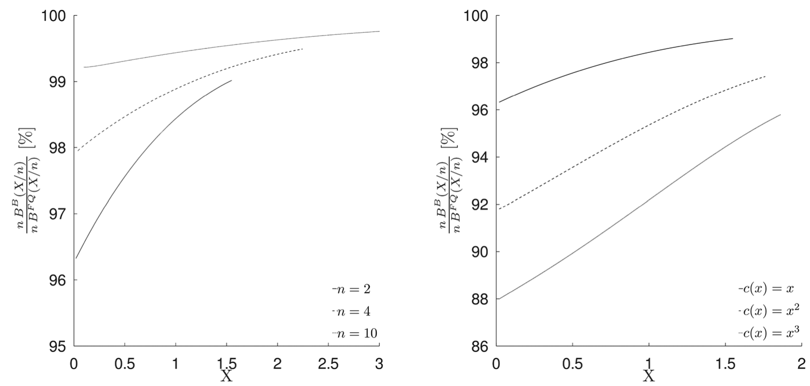

FIGURE 2 | Total costs of procurement in the budget auction as a share of total costs in the corresponding fixed-quantity auction. On the left: comparison for a budget auction with $c(x) = x$, $A \sim U[0, 1]$, and a differing number of players. On the right, the comparison is shown for a budget auction with $n = 2$, $A \sim U[0, 1]$, and differing cost structures.

5 | Limitations

5.1 | Uncertainty of the Budget

There are two reasons for introducing uncertainty regarding the budget. From a theoretical point of view, without uncertainty, the equilibrium from Proposition 2 only persists under the assumption that the bidders bid below their costs on units they do not supply in equilibrium. As this is weakly dominated by an otherwise unchanged strategy with above-cost bids for off-equilibrium quantities, such an equilibrium does not seem plausible in real auctions. As the distribution converges to a point mass on $\bar{A}$, bids converge to flat bids at level $c(\bar{x})$. The reason is that there is no cost of making the own bids on the infra-marginal units the bidder might lose, less attractive, as the density attached to these units is zero. However, the bidders have an incentive to deviate upward in this case. Consider the incentive of a bidder i to marginally raise their constant bids from $b_i = c(\bar{x}_i)$. Their profits for a constant $b_i \geq c(\bar{x}_i)$ are $\pi_i = A - \sum_{j \neq i} B_j(\beta_j(b_i)) - C(\frac{A - \sum_{j \neq i} B_j(\beta_j(b_i))}{b_i})$, where $A - \sum_{j \neq i} B_j(\beta_j(c(\bar{x}_i))) \equiv A_i$. If the other bidders bid weakly above costs for $x > \bar{x}_i$, the deviating bidder's profits are differentiable for a marginal upward deviation in b_i. Then, the benefit of marginally deviating upward is

$$\frac{\partial \pi_i}{\partial b_i}\Big|_{b_i = c(A_i/b_i)} = \frac{\partial A_i}{\partial b_i} - c\left(\frac{A_i}{b_i}\right)\left(\frac{\partial A_i}{\partial b_i}\frac{1}{b_i} - \frac{A_i}{(b_i)^2}\right) = \frac{A_i}{b_i} > 0.$$

Intuitively, the bidder only loses a non-profitable marginal unit but gains a margin on all other units. It is easy to verify that the same holds for fixed-quantity auctions.

From a practical point of view, some literature suggests that clear focal points make collusive behavior more likely (see, for example, Knittel and Stango 2003; Cramton and Ockenfels 2017). This is often a major concern of practitioners. An unknown budget can prevent collusion by eliminating the clear focal point of evenly splitting the budget. However, if procurement agencies are generally unable or unwilling to create uncertainty over the budget, the results of this article remain primarily theoretical.

5.2 | Uncertainty of the Bidders' Cost

The importance of an uncertain budget for the model hinges on the assumption of common knowledge of costs among the bidders. If their costs are private information, a symmetric equilibrium in pure strategies likely persists with a known budget, and collusion is harder as well. Consequently, an uncertain budget may not be necessary. Not considering private costs is the main limitation of the model in this article. However, there is reason to believe that this does alter the strategic effect of budget caps.

The budget effect probably carries over to a model with private costs and a known budget: A bidder raising their bids curbs the total demanded quantity and, thus, the quantity the bidder can win. Fundamentally, the bidders' problems share the same structure in both models, whether the probability of winning a certain number of units corresponds to the auctioneer's budget being sufficiently high or the other bidders not being too competitive.

5.3 | Optimal Mechanism with Individual Bidder Types

A model with individual bidder types and a known budget also allows for the analysis of the optimal mechanism for procuring the highest possible quantity with a given budget in a natural setup. Such a perspective complements the findings of this article. First, this article is not concerned with the *optimal* mechanism but with the strategic differences between budget and fixed-quantity auctions. Second, it does not model the auctioneer's uncertainty over the bidders' costs explicitly. Thus, it does not allow meaningfully answering the question of optimality.

If there is a continuum of bidder types, with high types having unambiguously higher costs than low types, and it is common knowledge that types represent independent draws from some distribution, the optimal mechanism for maximizing the procured quantity with a given budget can be derived analogously to Dasgupta and Spulber (1989). Participation constraints and incentive compatibility have standard implications, with no information rents accruing to the least efficient type and information rents reflecting the cost advantage of more efficient types. However, the auctioneer always spends the same amount, that is, the entire budget, as this allows them to procure a higher quantity. Consequently, when all bidders are of the least efficient type, demand is not distorted: as would be the case under perfect information, the auctioneer only pays for production costs and spends the entire budget. Thus, the same quantity as under perfect information is procured. By contrast, the bidders always make profits on inframarginal units in the budget auction.

5.4 | Endogenous Quality

The model also abstracts from quality considerations. These could weaken the result of Proposition 4 if the more competitive environment of the budget auction entailed a stronger incentive for bidders to lower product quality and costs. However, the marginal increase in expected profits a bidder can realize by lowering their costs is identical in both auction designs due to the dynamic equivalent of the standard envelope theorem (see Caputo 1990). As profits are maximized with respect to bids, a decrease in costs does not affect profits through bids. It only increases profits to the degree that it reduces costs for the (unchanged) equilibrium quantity. Thus, with the normalization of demand in Proposition 4, reducing costs is equally attractive under both designs.

5.5 | Preferences of the Auctioneer

Lower unit costs in the budget auction do not imply that the auctioneer necessarily prefers the budget auction. The article of Weitzman (1974) gives insights into how the preferences of the auctioneer, which are not modeled explicitly in this article, matter. In his model, the state either fixes a level of emissions or their price. The market outcome regarding the other variable is unknown to the state when it decides on its policy, as the costs of market participants are uncertain. In the model of Weitzman (1974), market participants do not behave strategically. Consequently, the costs of reaching a given level of emissions are identical under both policy options. Nonetheless, the state normally prefers one option: By fixing the level of emissions or their price for all realizations of costs, the state generally misses the optimal outcome in any given situation. Roughly speaking, the state should fix the price if the optimal price of emissions is almost identical for all cost realization, whereas the optimal level of emissions varies strongly and vice versa. An auctioneer who chooses between a budget and a fixed-quantity auction faces a similar problem. They fix either total spending or the traded quantity, and the not-fixed variable varies with the bidders' unknown costs. The insights of Weitzman (1974) imply that an auctioneer who wants to spend roughly the same amount (procure the same quantity) for any realization of the bidders' cost by tendency prefers a budget (fixed-quantity) auction. For example, an auctioneer who intrinsically values spending exactly their budget (for example, because of internal budgeting) tends to prefer a budget auction. In contrast to the model of Weitzman (1974), an auctioneer choosing between a budget and a fixed-quantity auction must also account for the strategic differences between the two. As the budget auction allows procuring a given quantity at lower total costs, it gains in attractiveness compared to a setting without strategic effects for most plausible preferences of the auctioneer.

6 | Conclusion

This article has considered a pay-as-bid procurement auction, in which the auctioneer maximizes the number of procured units given a secret budget constraint. In line with previous literature, the analysis assumed that the bidders are symmetric and share a common belief concerning the distribution of the auctioneer's budget. Deriving the symmetric equilibrium of the budget auction revealed that the budget auction introduces additional strategic complexity in comparison to the case of an uncertain but fixed quantity. In the budget auction, bidders not only weigh higher profit margins against a lower probability of realizing these same margins; higher bids also negatively influence the probability of winning later units. If a bidder demands a higher payment for some quantity, this reduces the auctioneer's residual budget such that it is less likely that the auctioneer's budget suffices to procure higher quantities from the bidder, which lowers bids. Thus, the budget cap leads to a link between bids in the bidders' problem, which has no correspondence in auctions with a fixed traded quantity. Due to this effect, bids in the budget auction are lower than in an auction with uncertain quantity, in which bidders believe in the same distribution of procured units as in the equilibrium of the budget auction.

Acknowledgments

I am grateful to Stefan Seifert for sparking my interest in the topic of this article and providing invaluable feedback. I am also deeply indebted to Fabian Herweg, Alessandro Kadner-Graziano, Christoph Kretschmer, three anonymous referees, and the editor Alessandro Bonatti for their extremely helpful comments and suggestions on earlier versions of this work. All remaining errors and shortcomings are entirely my own responsibility.

Open access funding enabled and organized by Projekt DEAL.

Endnotes

[1] Other possible solutions have been explored, for example, by Kremer and Nyborg (2004), Damianov (2005), and Burkett and Woodward (2020a).

[2] It only makes sense for the auctioneer to conduct a multi-unit auction, that is, source from multiple bidders, if marginal costs are increasing on some relevant interval.

[3] This is important because in most cases, the symmetric equilibrium in the budget auction can only be obtained numerically, which is impossible with an infinite slope.

[4] Due to both technical reasons and the fact that Holmberg (2009) considers electricity markets, he analyzes a somewhat different case, where demand exceeds production capacities with a positive probability. The model of Pycia and Woodward (2023) needs to be adapted to the case of a procurement auction.

[5] Technically, this formulation is somewhat imprecise, as a *single* bid does not influence a bidder's cumulative bids and, therefore, does not influence the probability of winning $y > x$ units.

[6] Normally, the Euler equation is a second-order differential equation. However, simultaneously determining all bidders' mutual best responses adds a degree of freedom.

[7] Off-equilibrium quantities are only relevant, if a deviating player reduces the maximum quantity they win. This would either require $b_i(\bar{x}_i) > c(\bar{x}_i)$, or $b_i(\bar{x}_i) = c(\bar{x}_i) < b_j(\bar{x}_j) = c(\bar{x}_j)$, where $i \neq j$ and $\bar{x}_j > \bar{x}_i$. In either case, the bidder could profitably win further units.

[8] The convenient property that the relevant part of the bidding function spans the same interval also explains the choice of the budget auction as the baseline for the comparison.

[9] Bids in the corresponding fixed-quantity auction are $b^{FQ}(x) = \frac{1}{2}x - \frac{\sqrt{15}}{5} + \frac{27\sqrt{3}\arccos(\frac{\sqrt{15}}{45}(5x+2\sqrt{15}))}{2\sqrt{225-75x^2-60\sqrt{15}x}}$.

[10] This is due to increased competition and total social costs for producing a given quantity decreasing with more producers.

[11] Stronger diseconomies of scale here are interpreted as "'more convex" cost functions. In the example, this is formalized analogously to utility functions with constant relative risk aversion. The finding regarding bidders' cost is not driven by the more convex cost functions resulting in lower costs for relevant quantities.

[12] The strategy $\hat{b}_i(x)$ is of course inadmissible as a solution because it has jump discontinuities at x_2 and x_3. However, the *optimal solution* is continuous, as can easily be seen from setting the bidder's problem up as a problem of optimal control, which would allow jump discontinuities in $b_i(x)$.

[13] Note that this proof is out of proper logical order.

References

Allen, J., R. Clark, B. Hickman, and E. Richert. 2024. "Resolving Failed Banks: Uncertainty, Multiple Bidding and Auction Design." *Review of Economic Studies* 91: 1201–1242.

Back, K., and J. F. Zender. 2001. "Auctions of Divisible Goods with Endogenous Supply." *Economics Letters* 73: 29–34.

Bergemann, D., and J. Hörner. 2018. "Should First-Price Auctions Be Transparent?" *American Economic Journal: Microeconomics* 10: 177–218.

Burkett, J., and K. Woodward. 2020a. "Reserve Prices Eliminate Low Revenue Equilibria in Uniform Price Auctions." *Games and Economic Behavior* 121: 297–306.

Burkett, J., and K. Woodward. 2020b. "Uniform Price Auctions with a Last Accepted Bid Pricing Rule." *Journal of Economic Theory* 185: 104954.

Caputo, M. R. 1990. "Comparative Dynamics via Envelope Methods in Variational Calculus." *The Review of Economic Studies* 57: 689–697.

Caputo, M. R. 2005. *Foundations of Dynamic Economic Analysis: Optimal Control Theory and Applications* Cambridge University Press.

Chiang, A. C. 1992. *Elements of Dynamic Optimization* Waveland Press.

Cramton, P., and A. Ockenfels. 2017. "The German 4G Spectrum Auction: Design and Behaviour." *The Economic Journal* 127: F305–F324.

Damianov, D. S. 2005. "The Uniform Price Auction with Endogenous Supply." *Economics Letters* 88: 152–158.

Dasgupta, S., and D. F. Spulber. 1989. "Managing Procurement Auctions." *Information Economics and Policy* 4: 5–29.

Dastidar, K. G. 2008. "On Procurement Auctions with Fixed Budgets." *Research in Economics* 62: 72–91.

Deck, C. A., and B. J. Wilson. 2008. "Fixed Revenue Auctions: Theory and Behavior." *Economic Inquiry* 46: 342–354.

Hansen, R. G. 1988. "Auctions with Endogenous Quantity." *The RAND Journal of Economics* 19: 44–58.

Holmberg, P. 2009. "Supply Function Equilibria of Pay-as-bid Auctions." *Journal of Regulatory Economics* 36: 154–177.

Klemperer, P. D., and M. A. Meyer. 1989. "Supply Function Equilibria in Oligopoly under Uncertainty." *Econometrica* 57: 1243–1277.

Knittel, C. R., and V. Stango. 2003. "Price Ceilings as Focal Points for Tacit Collusion: Evidence from Credit Cards." *American Economic Review* 93: 1703–1729.

Kremer, I., and K. G. Nyborg. 2004. "Underpricing and Market Power in Uniform Price Auctions." *The Review of Financial Studies* 17: 849–877.

Lengwiler, Y. 1999. "The Multiple Unit Auction with Variable Supply." *Economic Theory* 14: 373–392.

Liu, T., and C. Parlour. 2014. "Raising Money." Unpublished.

McAdams, D. 2007. "Adjustable Supply in Uniform Price Auctions: Non-Commitment as a Strategic Tool." *Economics Letters* 95: 48–53.

Oniki, H. 1973. "Comparative Dynamics (Sensitivity Analysis) in Optimal Control Theory." *Journal of Economic Theory* 6: 265–283.

Pycia, M., and K. Woodward. 2023. "Auctions of Homogeneous Goods: A Case for Pay-as-Bid." Working Paper. Accessed September 14, 2023. https://sites.econ.uzh.ch/pycia/pycia-woodward-paba.pdf.

Rothkopf, M. H., T. J. Teisberg, and E. P. Kahn. 1990. "Why Are Vickrey Auctions Rare?" *Journal of Political Economy* 98: 94–109.

RTR. 2020. "5G-Auktion macht nahezu flächendeckende Mobilfunk-Breitbandversorgung in Zukunft möglich." Accessed November 21, 2022. https://www.rtr.at/TKP/presse/pressemitteilungen/pressemitteilungen/pinfo11092020tkp.de.html.

Tillio, A. d., N. Kos, and M. Messner. 2016. "The Design of Ambiguous Mechanisms." *The Review of Economic Studies* 84: 237–276.

TKK. 2019. "Tender Document in the Procedure for Awarding Spectrum in the 700, 1500 and 2100 MHz Ranges." Accessed December 16, 2022. https://www.rtr.at/TKP/was_wir_tun/telekommunikation/spectrum/procedures/Multibandauktion_700-1500-2100MHz_2020/FRQ5G_2020_tender_document.de.html.

Weitzman, M. L. 1974. "Prices vs. Quantities." *The Review of Economic Studies* 41: 477–491.

Wilson, R. 1979. "Auctions of Shares." *The Quarterly Journal of Economics* 93: 675–689.

Appendix A: Proof of Proposition 2

A.1 | First-Order Condition

The relevant derivatives for the first-order condition given in (5) are

$$\frac{\partial \Psi}{\partial B_i} = -[b_i(x) - c(x)] f(B_i(x) + \sum_{j \neq i} B_j(\beta_j(b_i(x)))), \tag{A1}$$

$$\begin{aligned} \tfrac{\partial \Psi}{\partial b_i} = & \quad 1 - F(B_i(x) + \textstyle\sum_{j \neq i} B_j(\beta_j(b_i(x)))) \\ & -[b_i(x) - c(x)] f(B_i(x) + \textstyle\sum_{j \neq i} B_j(\beta_j(b_i(x)))) \textstyle\sum_{j \neq i} b_i(x) \beta_j'(b_i(x)), \end{aligned} \tag{A2}$$

and

$$\begin{aligned} \frac{d}{dx}\left(\frac{\partial \Psi}{\partial b_i}\right) = & -f(\cdot) \left[b_i(x) + \sum_{j \neq i} b_i(x)\, \beta_j'(b_i(x))\, b_i'(x) \right] \\ & - \left[b_i'(x) - c'(x) \right] f(\cdot) \sum_{j \neq i} b_i(x)\, \beta_j'(b_i(x)) \\ & -[b_i(x) - c(x)] \left\{ f'(\cdot) \left[b_i(x) + \sum_{j \neq i} b_i(x)\, \beta_j'(b_i(x))\, b_i'(x) \right] \right. \\ & \times \sum_{j \neq i} b_i(x)\, \beta_j'(b_i(x)) + f(\cdot) \sum_{j \neq i} b_i'(x)\, \beta_j'(b_i(x)) \\ & \left. + f(\cdot) \sum_{j \neq i} b_i(x)\, \beta_j''(b_i(x))\, b_i'(x) \right\}. \end{aligned} \tag{A3}$$

Searching for a symmetric equilibrium with $b(x) = b_i(x) = b_j(x) \forall j$, furthermore gives $\beta_j'(b_i(x)) = \frac{1}{b'(x)}$ and $\beta_j''(b_i(x)) = -\frac{b''(x)}{[b'(x)]^3}$. Using these symmetry conditions together with the derivatives in Equations (A1) and (A3) in the necessary condition from Equation (5) gives Equation (6).

A.2 | Definitizing the (numeric) Solution to the Characteristic Differential Equation

Definitizing $B(\bar{x})$ **and** $b(\bar{x})$: The highest quantity a given bidder can win, is determined according to the terminal curve $B_i(\bar{x}_i) = \phi(b_i, \bar{x}_i)$. Based on the general transversality condition for the endpoint of a variational problem, the Euler-Lagrange equation of the considered bidder's best response problem from Equation (5) is completed by $B_i(\bar{x}_i) = \phi(b_i, \bar{x}_i)$ and $[\Psi(\cdot) + (\frac{d}{d\bar{x}_i}\phi(\cdot) - b_i(x))\frac{\partial \Psi}{\partial b_i}]_{x=\bar{x}_i} = 0$ (see e.g., Chiang 1992). These conditions translate to

$$B_i(\bar{x}_i) + \sum_{j \neq i} B_j\left(\beta_j(b_i(\bar{x}_i))\right) = \bar{A} \tag{A4}$$

and given Equation (A4)

$$\left(\sum_{j \neq i} b_i(\bar{x}_i)\, \beta_j'(b_i(\bar{x}_i))\, b_i'(\bar{x}_i) + b_i(x) \right) [b_i(\bar{x}_i) - c(\bar{x}_i)]\, f(\bar{A}) \sum_{j \neq i} b_i(\bar{x}_i)\, \beta_j'(b_i(x)) = 0. \tag{A5}$$

By assumption $f(\bar{A}) > 0$ and bids of all players are positive and increasing, such that Equation (A5) implies that

$$b_i(\bar{x}_i) = c(\bar{x}_i). \tag{A6}$$

For the symmetric equilibrium, the transversality conditions from Equations (A4) and (A6) give Equations (7) and (8).

Definitizing $b'(\bar{x})$: Given that for a response of i to be optimal, it must be that $b_i(\bar{x}_i) = c(\bar{x}_i)$, the first-order condition of bidder i's problem obtained by using Equations (A1) and (A3) in Equation (5) simplifies substantially. Making use of the assumptions that $f(\bar{A})$, $f'(\bar{A})$, as well as the derivatives of the bidding functions, are finite, several terms vanish such that the simplified first-order condition emerges as

$$b_i'(\bar{x}_i) = \frac{c'(\bar{x}_i)}{2} - \frac{1}{2\sum_{j \neq i} \beta_j'(c(\bar{x}_i))}. \tag{A7}$$

As, in equilibrium, this condition must hold for every player, we can use symmetry to obtain the formulation in Equation (9).

Definitizing $b''(\bar{x})$: The formulation of $b''(\bar{x})$, as given in (6), cannot be used for the numerical approximation of the symmetric equilibrium, as it is indeterminate. This can be seen by substituting the expression for $b'(\bar{x})$ from (9) and noting that $b(\bar{x}) - c(\bar{x}) = 0$. Applying Hôpital's rule to (6) yields

$$b''(\bar{x}) = \frac{[b'(\bar{x})]^2}{b(\bar{x})\,[b'(\bar{x})-2c'(\bar{x})]}\Big[3b'(\bar{x}) - \frac{(2n-3)}{(n-1)}\,c'(\bar{x}) - c''(\bar{x})\,\frac{b(\bar{x})}{b'(\bar{x})} + n\,\frac{f'(nB(\bar{x}))}{f(nB(\bar{x}))}\,[b'(\bar{x}) - c'(\bar{x})]\,\frac{[b(\bar{x})]^2}{b'(\bar{x})}\Big]. \tag{A8}$$

A.3 | Properties of Equilibrium Bids

Bids are Above Marginal Costs: Consider a bidder i with bidding function $b_i(x)$ who, on some arbitrarily small, relevant interval $[x_1, x_2]$ of strictly positive length, bids weakly below marginal costs. Furthermore, assume that the bidder also has a weakly positive margin on the interval $[x_2, x_3]$ with $x_3 - x_2 = x_2 - x_1$ and $x_3 \le \bar{x}_i$. Consider the alternative bidding function

$$\hat{b}_i(x) = \begin{cases} b_i(x) + \varepsilon & \text{for } x \in [x_1, x_2] \\ b_i(x) - \delta\varepsilon & \text{for } x \in (x_2, x_3] \\ b_i(x) & \text{else,} \end{cases}$$

where $\delta \ge 1$ and $x_3 = [(1+\delta)x_2 - x_1)]/\delta$. For an appropriate choice of δ the bidder would profit from choosing a positive ε. Because both $b_i(x)$ and $B_i(x)$ are unchanged outside the interval $[x_1, x_3]$, the marginal effect on profits of increasing ε from 0 is

$$\begin{aligned} \frac{\partial \mathbb{E}[\pi_i]}{\partial \varepsilon}\Big|_{\varepsilon=0} = & \int_{x_1}^{x_2}\left[1 - F\left(B_i(x) + \sum_{j\neq i} B_j\left(\beta_j(b_i(x))\right)\right)\right]dx - \int_{x_2}^{x_3(\delta)} \delta[1 - F(\cdot)]\,dx \\ & - \int_{x_1}^{x_2} \underbrace{[b_i(x) - c(x)]\,f(\cdot)\left[(x - x_1) + \sum_{j\neq i} b_i(x)\beta_j'(b_i(x))\right]}_{\le 0} dx \\ & + \int_{x_2}^{x_3(\delta)} [b_i(x) - c(x)]\,f(\cdot)\,\delta \sum_{j\neq i} b_i(x)\beta_j'(b_i(x))\,dx \\ & - \int_{x_2}^{x_3(\delta)} [b_i(x) - c(x)]\,f(\cdot)\,[(x_2 - x_1) - \delta(x - x_2)]\,dx. \end{aligned} \tag{A9}$$

For high enough δ, $\frac{\partial \mathbb{E}[\pi_i]}{\partial \varepsilon}|_{\varepsilon=0} > 0$. To see this, note that the first three lines are each (weakly) positive. The first line is positive because $\int_{x_1}^{x_2} 1 - F(\cdot)dx > (x_2 - x_1)[1 - F(\cdot)]_{x=x_2} = \delta(x_3 - x_2)[1 - F(\cdot)]_{x=x_2} > \int_{x_2}^{x_3} \delta[1 - F(\cdot)]dx$. The second line is positive due to the weakly negative margins in the interval. Lastly, all terms in the integrand in the third line are positive. The integral in the fourth line is decreasing in δ and converges to zero as $\delta \to \infty$. It follows that for high enough δ, the expression in (A9) is positive and there is a profitable deviation from $b_i(x)$.[12] This argument also holds if the bidder does not bid above costs for $x > x_2$. As a consequence, the profitability of a deviation to $\hat{b}_i(x)$ shows that a bidder's best response cannot include below-cost bidding.

Monotonicity of Bids: Equation (9) establishes that $b'(\bar{x}) > 0$. Against the backdrop of Equation (B2) the slope of equilibrium bids for $x \in [0, \bar{x})$ can be written as[13]

$$b'(x) = b'_{FQ}(x) - \Delta b'(x). \tag{A10}$$

Given Equations (B1) and (B3), we can write

$$b'(x) = (n-1)\,[1 - F(\cdot)]^{-1}\,f(\cdot)\,b(x)\,[b(x) - c(x)] - (n-1)\,[1 - F(\cdot)]^{\frac{1-n}{n}}\,\Delta\tilde{b}'(x) > 0. \tag{A11}$$

It was shown above that $b(x) > c(x)$ for $x \in [0, \bar{x})$ and Appendix B shows that $\Delta\tilde{b}'(x) < 0$ for $x \in [0, \bar{x})$. Thus, Equation (A11) establishes the monotonicity of bids.

A.4 | Second-Order Condition

General Case: The most straightforward way to ensure that the symmetric equilibrium candidate described by Equations (6)–(9) does indeed constitute an equilibrium is to make sure that the integrand of a unilaterally deviating bidder's profit function is globally concave. Although, depending on the case at hand, weaker conditions might be given, I will state the conditions for the strict concavity of $\Psi(\cdot)$. For the deviating bidder i, the integrand of their expected profits is

$$\Psi(B_i, b_i, x) = [b_i(x) - c(x)]\,[1 - F(B_i(x) + (n-1)\,B(\beta(b_i(x)))]. \tag{A12}$$

Strict global concavity requires that for any B_i and b_i the corresponding quadratic form is negative definite. In particular, it must be the case that

$$\frac{\partial^2\Psi}{\partial(b_i)^2} < 0 \text{ and } \frac{\partial^2\Psi}{\partial(b_i)^2} \times \frac{\partial^2\Psi}{\partial(B_i)^2} - \left(\frac{\partial^2\Psi}{\partial b_i\,\partial B_i}\right)^2 < 0 \tag{A13}$$

for all (B_i, b_i), where

$$\frac{\partial^2\Psi}{\partial(b_i)^2} = -2(n-1)f(\cdot)b_i(x)\beta'(\cdot) - [b_i(x)-c(x)]f'(\cdot)\left[(n-1)b_i(x)\beta'(\cdot)\right]^2 - (n-1)[b_i(x)-c(x)]f(\cdot)\left(\beta'(\cdot)+b_i(x)\beta''(\cdot)\right),$$

$$\frac{\partial^2\Psi}{\partial(B_i)^2} = -[b_i(x)-c(x)]f'(\cdot) \text{ and } \frac{\partial^2\Psi}{\partial b_i\,\partial B_i} = -f(\cdot)-(n-1)[b_i(x)-c(x)]f'(\cdot)b_i(x)\beta'(\cdot).$$

The condition of global concavity can be weakened, as any maximum of a bidder's expected profits can only be located in the area where $b_i(x) > c(x)$.

Example from Section 3.3: Assume that the continuation of the bidding function for off-equilibrium quantities is differentiable and weakly concave. The problem satisfies concavity given that only cases with $b_i(x) - c(x)$ need to be considered. In this case $\frac{\partial^2\Psi}{\partial(b_i)^2} < 0$. Furthermore, due to the uniform distribution $\frac{\partial^2\Psi}{\partial(B_i)^2} = 0$ and $\frac{\partial^2\Psi}{\partial b_i\,\partial B_i} = -1$. Thus, the second-order condition is satisfied in the example.

Appendix B: Proof of Proposition 4

The main text establishes that $b^{FQ}(\bar{x}) = b^B(\bar{x})$. It remains to be shown that $b^{FQ}(x) > b^B(x)$ for $x \in [0, \bar{x})$. Based on Equations (3) and (10), bids in the corresponding fixed-quantity auction are

$$b^{FQ}(x) = (n-1)\left[1-F\left(nB^B(x)\right)\right]^{\frac{1-n}{n}}\int_x^{\bar{x}}[1-F(\cdot)]^{-1/n}f(\cdot)\,b^B(y)\,c(y)\,dy. \tag{B1}$$

Noting that bids in the budget auction can be written as

$$\begin{aligned} b^B(x) &= (n-1)\left[1-F\left(nB^B(x)\right)\right]^{\frac{1-n}{n}}\mu(x)\int_x^{\bar{x}}[1-F(\cdot)]^{-1/n}f(\cdot)\,b^B(y)\,b^B(y)\,dy, \\ \text{where } \mu(x) &= \frac{b^B(x)}{b^B(x)+\left[1-F\left(nB^B(x)\right)\right]^{\frac{1-n}{n}}\int_x^{\bar{x}}\left[1-F\left(nB^B(y)\right)\right]^{\frac{n-1}{n}}b^{B\prime}(y)\,dy}, \end{aligned}$$

the difference in bids, $\Delta b(x) = b^{FQ}(x) - b^B(x)$, is

$$\Delta b(x) = (n-1)[1-F(\cdot)]^{\frac{1-n}{n}}\int_x^{\bar{x}}[1-F(\cdot)]^{-1/n}f(\cdot)\,b^B(y)\,[c(y)-\mu(x)\,b^B(y)]\,dy. \tag{B2}$$

Given that we are considering $x < \bar{x}$, the hypothesis that $\Delta b(x) > 0$ is equivalent to

$$\Delta\tilde{b}(x) = \int_x^{\bar{x}}\left[1-F\left(nB^B(y)\right)\right]^{-1/n}f\left(nB^B(y)\right)b^B(y)\,[c(y)-\mu(x)\,b^B(y)]\,dy > 0. \tag{B3}$$

Because $\lim_{x\to\bar{x}}\Delta\tilde{b}(x) = 0$, a sufficient condition for inequality (B3) and by extension $\Delta b(x) > 0$ for $x < \bar{x}$ to hold is that $\Delta\tilde{b}'(x) < 0$, that is,

$$\Delta\tilde{b}'(x) = [1-F(\cdot)]^{-1/n}f(\cdot)\,b^B(x)\left[\mu(x)\,b^B(x)-c(x)\right] - \mu'(x)\int_x^{\bar{x}}[1-F(\cdot)]^{-1/n}f(\cdot)\,[b^B(y)]^2\,dy < 0. \tag{B4}$$

Some manipulation reveals that inequality (B4) is equivalent to

$$b^B(x) - c(x) - \frac{1}{(n-1)}\,\frac{1-F\left(nB^B(x)\right)}{f(nB^B(x))}\,\frac{b^{B\prime}(x)}{b^B(x)} < 0. \tag{B5}$$

A comparison with Equation (A2) shows that inequality (B5) is equivalent to $\frac{\partial\Psi}{\partial b_i}\Big|_{B_i=B_j=B^B,b_i=b_j=b^B} > 0$. Given that bids are above costs (see Appendix A.3), Equation (A1) reveals that $\frac{\partial\Psi}{\partial B_i} < 0$ for $x \in [0, \bar{x})$ if evaluated along the equilibrium path. Thus, the first-order condition from (5) implies that $\frac{d}{dx}\frac{\partial\Psi}{\partial b_i} < 0$ for $x < \bar{x}$ in equilibrium. Using the terminal conditions from Equations (7) and (8) in (A2) gives that, evaluated at equilibrium bids, $\frac{\partial\Psi}{\partial b_i}\Big|_{x=\bar{x}} = 0$. It follows that $\frac{\partial\Psi}{\partial b_i}\Big|_{B_i=B_j=B^B,b_i=b_j=b^B} > 0$, that is, inequality (B5) holds.

B
Submitted Version of Chapter 5

Watt's Better?
Pay-as-Bid vs. Uniform Pricing in Electricity Markets

Claudio Rottner[*][†]

June 27, 2025

Abstract

This paper derives the symmetric equilibrium in an oligopolistic electricity market where symmetric producers – each with a weakly increasing marginal cost function – compete for an uncertain, downward-sloping demand under pay-as-bid pricing. Comparing pay-as-bid and uniform pricing, this paper finds that both price rules implement first-best welfare with zero profits when marginal costs are flat or there are infinitely many producers. For a linear demand function with a uniformly distributed random intercept and linearly increasing marginal costs, pay-as-bid pricing results in a higher expected consumer surplus. Pay-as-bid pricing results in higher welfare in high-demand situations, whereas either pricing rule can perform better for low demand. Thus, the general comparison of expected welfare remains ambiguous in the linear model. However, lower demand variability favours pay-as-bid pricing; for low enough demand variability pay-as-bid pricing results in higher expected welfare; if demand variability is very low, pay-as-bid pricing even results in higher welfare for every demand realisation. These findings suggest that regulators should consider pay-as-bid pricing to increase consumer surplus. Such a switch is especially promising in short-term markets with low demand uncertainty, where a positive effect on welfare is most likely.

Keywords: Electricity Markets, Auctions, Uniform Price, Pay-as-Bid

JEL Codes: C72, D44, D82, L13

[*]Corresponding author; University of Bayreuth, Universitätsstraße 30, 95447 Bayreuth, Germany; e-mail claudio.rottner@uni-bayreuth.de

[†]This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

1 Introduction

On January 8th, 2025, the price of a MWh of electricity on the intraday market for Great Britain peaked at more than £1,000/MWh (Ambrose, 2025) – roughly ten times the going rate for the whole month (EPEX, 2025), and several times the estimated production costs for even the most expensive technologies (International Energy Agency, 2020; Department for Energy Security & Net Zero, 2023a). Such price peaks raise questions about the adequacy of the current market institutions, particularly, the way prices are set in electricity markets.

Currently, under the uniform-price rule, every unit is priced at a the same price determined by the most expensive unit traded. Intuitively, this puts the burden of high-price situations disproportionately on consumers, who pay high prices even on those units that are cheap to produce. Consequently, it is unsurprising that this price rule has come under fire with the rising energy prices of recent years. For example, in her State of the Union address in 2022, Ursula von der Leyen, president of the European Commission, argued that "[t]he current electricity market design – based on merit order – is not doing justice to consumers anymore. They should reap the benefits of low-cost renewables. So, we have to decouple the dominant influence of gas [usually the most expensive technology] on the price of electricity." (European Commission, 2022). An obvious way to do so would be to change the pricing rule to what is called pay-as-bid, or discriminatory, pricing, where every unit trades at "its own" price, which corresponds to the ask-price of the producer. For example, the UK government's Review of Electricity Market Arrangements, which started in 2022, considered this switch. However, the idea was dismissed early in the consultations: Its proponents argued that it would benefit consumers; by contrast, its opponents warned of "the risk of tactical bidding" which could lead to lower efficiency (Department for Energy Security & Net Zero, 2023b).

These diverging assessments reflect that the literature has failed to analytically describe bidding behaviour under the pay-as-bid rule in conditions plausibly reflecting electricity markets. Whereas equilibrium behaviour under the uniform pricing rule, or equivalently, in uniform-price auctions, has been well understood going back to Klemperer and Meyer (1989), equilibrium in pay-as-bid auctions has mostly eluded a proper analytic solution. Therefore, comparisons of the two pricing rules have relied on analytic models with strong simplifying assumptions (e.g., Federico and Rahman, 2003; Fabra et al., 2006; Hästö and Holmberg, 2006; Genc, 2008; Holmberg, 2008, 2009) or on simulations (e.g., Bower and Bunn, 2001; Rassenti et al., 2003; Xiong et al., 2004; Guerci et al., 2007; Sugianto and Liao, 2014; Anatolitis and Welisch, 2017; Viehmann et al., 2021) which make it hard to identify the relevant factors. However, recent advances in auction theory (especially Pycia and Woodward, 2025) make it possible to obtain wider analytic results in comparing the two pricing rules on electricity markets. This paper takes a step in this direction.

The contribution of this paper is twofold: First, it derives bidding behaviour in the form of pure-strategy supply function equilibria under the pay-as-bid rule for an oligopoly with a general, random, downward-sloping demand for arbitrary, increasing marginal costs. Furthermore, this paper derives some basic insights regarding the comparison of the pricing rules in this general setup. Second, it develops richer results regarding the two pricing rules in a linear model where marginal costs are linearly increasing and market demand is linear with a uniformly-distributed intercept.

The general model shows that both pricing rules implement the first-best welfare with zero profits for producers when marginal costs are constant (as much of the previous literature assumes) or there are infinitely many suppliers. Otherwise, the comparison of consumer surplus and expected welfare is, in general, ambiguous. In the linear version of the model, pay-as-bid pricing results in a higher expected consumer surplus because it lowers prices on inframarginal units. The comparison of expected welfare remains ambiguous: Pay-as-bid pricing realises higher welfare in high demand situations, whereas either rule can do better when demand is low. If demand uncertainty (measured by the geometric variance of the intercept) is low, pay-as-bid pricing results in higher expected welfare. If demand uncertainty is very low, it even outperforms uniform pricing for every demand realisation.

The results of this paper suggest that consumers would benefit from a switch to pay-as-bid pricing. This change in the pricing rule is most likely to increase overall welfare when demand uncertainty is low. Thus, pay-as-bid pricing is more promising in short-term markets (e.g., the spot market) than in longer-term markets. Furthermore, this paper shows that optimal bidding behaviour under pay-as-bid pricing in electricity markets can be described under fairly general conditions; fears of erratic bidding behaviour are unfounded.

This paper proceeds as follows. Section 2 discusses related literature. Section 3 introduces the model and derives equilibrium behaviour before Section 4 compares the performance of the two pricing rules regarding consumer surplus and welfare. Section 5 discusses the results of this paper and their implications. Section 6 concludes.

2 Related Literature

The first contribution of this paper is to derive the symmetric equilibrium under pay-as-bid pricing under assumptions plausibly reflecting electricity markets: Supply is oligopolistic, demand is variable and downward-sloping (albeit inelastic), and marginal costs are increasing. Previous models of pay-as-bid pricing have neglected at least one of these aspects.[1]

Some models, for example, Fabra et al. (2006) and Genc (2008), assume constant unit costs. This assumption is unrealistic given different production technologies and impacts results in two important ways: First, if producers are symmetric and not capacity constrained, the pay-as-bid auction collapses into Bertrand competition with truthful bidding (see Genc, 2008); producers can only make margins to the degree that the competitors cannot undercut this price because they are in a worse cost position. With constant marginal cost, truthful bidding is also an equilibrium in the uniform-price auction. Consequently, with flat marginal costs, both pricing rules implement the first-best welfare with zero profits for producers (see Corollary 4); this seems a hugely inadequate description of real-world electricity markets. Second, even when constant unit costs do not imply perfect competition because produc-

[1] Outside the context of electricity markets, some authors also restrict attention to single-unit demand/supply of bidders (e.g., Krishna, 2010; Anderson and Holmberg, 2023; Bougt et al., 2025).

ers are asymmetric or capacity-constrained, as in Fabra et al. (2006), this modeling choice renders mute the most common argument against uniform pricing: Consumer prices on inframarginal units are not excessive in a uniform-price auction because inframarginal units come at the same cost to producers as marginal units.

Another common assumption is completely inelastic demand (e.g., Genc, 2008; Hästö and Holmberg, 2006; Holmberg, 2008, 2009; Pycia and Woodward, 2025). However, even in spot markets, electricity demand is not completely inelastic. Assuming otherwise can dramatically change results. In the model of this paper, prices in a uniform-price auction can go to infinity if they are not bounded by consumers' willingness to pay.[2]

Last, some papers ignore the oligopolistic structure of electricity markets. For example, Federico and Rahman (2003) consider the equivalent of the linear version of the model in this paper. However, they restrict attention to the limiting cases of perfect competition and monopoly. Consequently, their treatment of equilibrium behaviour under pay-as-bid pricing falls short of the analysis in this paper regarding the market structure and the generality of the assumptions regarding costs and demand.

Regarding the comparison of the pricing rules, this paper makes the same functional form assumptions as Federico and Rahman (2003) for the most part. Thus, these comparisons essentially generalise the results of Federico and Rahman (2003) to the oligopolistic setting. Their results are broadly consistent with this paper. They too find that the expected consumer surplus is higher under pay-as-bid pricing. For expected welfare, they find that the uniform-price auction performs better under perfect competition. By contrast, welfare converges to first-best welfare under both pricing rules in the analogue of perfect competition, i.e., when there are infinitely many producers, in the model of this paper (see the discussion of Corollary 3 for an explanation). In the monopoly case, they too find that neither price rule generally results in higher welfare. Consistent with this paper, they show that flat marginal costs and low demand uncertainty favour the pay-as-bid auction.

[2]This is not specific to this model. Genc (2008) and Holmberg (2008, 2009) circumvent this problem by assuming that prices are capped.

Apart from these analytic models, most of the literature comparing uniform and pay-as-bid pricing in electricity markets relies on simulations to derive the equilibrium strategies of producers (e.g., Rassenti et al., 2003; Xiong et al., 2004; Guerci et al., 2007; Sugianto and Liao, 2014; Anatolitis and Welisch, 2017; Viehmann et al., 2021). On the one hand, this approach makes it hard to generalise the results given that the knowledge of what drives them is limited. On the other hand, simulations allow for considering more complex scenarios. For example, many simulations consider the repeated nature of these auctions and the resulting learning (e.g., Bower and Bunn, 2001; Anatolitis and Welisch, 2017; Viehmann et al., 2021), whereas the analytic literature, including this paper, models a simple one-off auction. Moreover, Bower and Bunn (2001) and Rassenti et al. (2003) also consider the role of asymmetry in the size of producers. The results of these simulation studies are mixed. Xiong et al. (2004), Anatolitis and Welisch (2017), and Viehmann et al. (2021) seem to suggest that consumer surplus is higher under pay-as-bid pricing, whereas Bower and Bunn (2001) and Rassenti et al. (2003) support the opposite conclusion. Similarly, some studies (e.g., Anatolitis and Welisch, 2017; Viehmann et al., 2021) favour uniform pricing regarding welfare, while others (e.g., Guerci et al., 2007) do not.

3 Model and Bidding Behaviour

3.1 Model

Demand.— The demanded quantity in the entire market at price p is given by $D(p, \varepsilon)$, where ε is a demand-shifter. Without loss of generality, assume that $\frac{\partial D(p,\varepsilon)}{\partial \varepsilon} > 0$. The demand-shifter ε is a random variable that follows follows the cumulative distribution $F(\varepsilon)$ on $[\underline{\varepsilon}; \bar{\varepsilon}]$ with $F'(\varepsilon) = f(\varepsilon) > 0$. This paper assumes that demand is downward-sloping, i.e., that $\frac{\partial D(p,\varepsilon)}{\partial p} < 0$. However, the case of inelastic demand is nested in the derivation of equilibrium behaviour. Throughout the paper, assume that $D(p, \varepsilon)$ reflects consumers' preferences without strategic distortion.

Supply.— There are $n \geq 2$ symmetric producers that compete in the market. Each producer has the same, commonly known marginal cost function $c(x)$ with $c'(x) \geq 0$.[3] Every bidder i submits a weakly increasing bidding function $\beta_i(x)$ that reflects, how much money they demand when supplying x units.

Market Clearing.— The market clears at a stop-out price p^* that equalises demand and aggregate supply, i.e.,

$$D(p^*, \varepsilon) = \sum_i \beta_i^{-1}(p^*), \tag{1}$$

where $\beta_i^{-1}(p)$ is the inverse of $\beta_i(x)$. Thus, as long as bidding functions are strictly increasing, a producer supplies $\beta_i^{-1}(p^*)$ units, i.e., all units for which their bidding function is weakly below p^*. Throughout the paper, assume that the minimal demand level is high enough for trade to occur for every demand realisation under both pricing rules.

Pricing Rules.— Under the uniform-price rule, a buyer is paid the stop-out price p^* for all units they supply. When they bid $\beta_i(x)$ and the stop-out price is p^*, they make profits of

$$\pi_{UP} = p^* \; \beta_i^{-1}(p^*) - C\big(\beta_i^{-1}(p^*)\big),$$

where $C(x) = \int_0^x c(y)\,dy$. Under the pay-as-bid rule, they are paid the price specified by their bidding function $\beta_i(x)$ for each supplied unit; their profits with stop-out price p^* are

$$\pi_{PAB} = \int_0^{\beta_i^{-1}(p^*)} \beta_i(x)\,dx - C\big(\beta_i^{-1}(p^*)\big).$$

Figure 1 illustrates a producer's revenue under the two pricing rules.

Discussion of Model Assumptions.— Even though the fundamental model assumptions are typical in the context of electricity markets, they depart (with good reason) from some of the assumptions traditionally made in auction theory.

First, demand is exogenous and not strategically determined by the auctioneer. As

[3] Equilibrium bidding strategies do not change when including fixed costs.

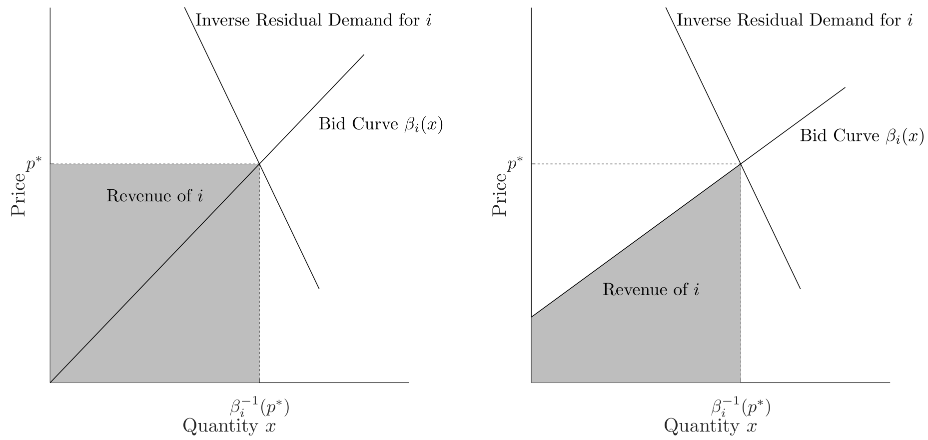

(a) Uniform pricing. (b) Pay-as-bid pricing.

Figure 1: Illustration of uniform and pay-as-bid pricing.

regulators are not active in the market, they cannot change demand, which is, therefore, exogenous to their decision. Even though the demand side does not determine the pricing rule, it presumably reacts strategically to it (see Kamat and Oren, 2002, for such a model); this paper neglects strategic demand. However, the model assumes that demand varies, capturing (uncertainty over) time-varying demand.

Second, the producers' costs are not private information (see Burkett and Woodward, 2020, for a model of a uniform price auction with private information) but common knowledge. This assumption seems plausible for electricity markets given that a producer's power plants are generally observable to its competitors, who use the same production technologies. The model abstracts from the fact that in electricity markets production costs too vary over time. Given that none of the results of this paper rely on the level of costs, the results generalise to varying cost levels as long as the producers' uncertainty over their own and their competitors' costs resolves before bidding takes place.

Linear Model.— For the case of constant marginal costs or infinitely many suppliers, Corollaries 3 and 4 show that both pricing rules implement first-best welfare with all welfare going to consumers. All other comparisons rely on a linear model, where

$$D(p, \varepsilon) = \varepsilon - mp, \tag{2a}$$

$$\varepsilon \sim U[\underline{\varepsilon}; \bar{\varepsilon}], \tag{2b}$$

and

$$c(x) = \alpha x \tag{2c}$$

with $m, \alpha, \underline{\varepsilon} > 0$ and $\bar{\varepsilon} > \underline{\varepsilon} > 0$.

3.2 Bidding Behaviour under Pay-as-Bid Pricing

This section derives equilibrium bidding behaviour under pay-as-bid pricing in the general model with an arbitrary demand and cost function, and for the linear version of the model.

Bidder's Problem.— Under pay-as-bid pricing, a producer i earns their margin $\beta_i(x) - c(x)$ on some unit x if and only if they get to supply weakly more than x units. Assume all other producers $j \neq i$ follow the strategy $\beta_j(x)$. In that case, producer i supplies weakly more than x units, if and only if demand at price $\beta_i(x)$ is large enough to cover the x units of producer i and the $(n-1)\,\beta_j^{-1}(\beta_i(x))$ units the other producers supply at this price. Thus, producer i earns the margin $\beta_i(x) - c(x)$ if and only if $D(\beta_i(x), \varepsilon) \geq x + (n-1)\beta_j^{-1}(\beta_i(x))$. This is equivalent to the demand-shifter ε taking a high enough value, specifically,

$$\varepsilon \geq e\Big(b_i(x), x + (n-1)\beta^{-1}(b_i(x))\Big),$$

where $e(p, Q)$ it the unique value of ε for which $D(p, e) = Q$, i.e., $D\big(p, e(p, Q)\big) = Q$. Thus, bidder i's equilibrium bidding strategy solves the maximisation problem

$$\max_{\beta_i(x)} \qquad \int_{\underline{x}_i}^{\bar{x}_i} [\beta_i(x) - c(x)] \Big[1 - F\Big(e\big(\beta_i(x), x + (n-1)\beta_j^{-1}(\beta_i(x))\big)\Big)\Big] \, dx, \tag{3}$$

where $\underline{x}_i$ and $\bar{x}_i$ denote the quantities bidder i gets to supply with their strategy when demand is at its minimal level $\underline{\varepsilon}$ and maximum level $\bar{\varepsilon}$, respectively. The objective function in (3) corresponds to the margin on unit x, weighted by the probability that the demand-shifter ε is large enough for the producer to supply more than x units.

The maximisation problem in (3) can be solved by point-wise maximisation of the integrand; the first-order condition is

$$1 - F(\cdot) = [\beta_i(x) - c(x)]\, f(\cdot)\Big(e_p(\cdot) + e_Q(\cdot)(n-1)\beta_j^{-1\prime}(\beta_i(x))\Big). \tag{4}$$

The producer benefits from increasing their bid on some unit x by making additional profits whenever unit x is inframarginal, i.e., whenever they win more than x units (left-hand side); on the downside, they lose the margin on the unit in the case where they do not get to supply the unit because of the higher bid, i.e., when the unit is marginal (right-hand side). As the unit $\bar{x}_i$, which the producer only supplies when $\varepsilon = \bar{\varepsilon}$, is never inframarginal, there is no incentive to bid above marginal costs on that unit. Consequently, producer i's optimal bidding function satisfies

$$\beta_i(\bar{x}_i) = c(\bar{x}_i).^4 \tag{5}$$

The trade-off in the first-order condition from (4) only captures the producer's problem correctly, as long as they are at risk of losing the unit in question: In that case, the producer trades off a positive probability of realising the margin on that unit (positive left-hand side) with the positive probability of losing that margin if demand is sufficiently low (positive right-hand side). However, for units $x < \underline{x}_i$, there is no risk of losing them, as the producer supplies $\underline{x}_i$ units even if demand is at its minimal level. Evidently, the buyer should increase the bid on these units until they become marginal for the minimal demand level; thus, the

[4]Formally, the left-hand side of (4) is zero for $\varepsilon = \bar{\varepsilon}$. Thus, the right-hand side must be zero to satisfy (4). Given that with increasing bids all but the first factor are positive, (5) must hold.

bid on all units $x \leq \underline{x}_i$ is identical. Thus, the optimal bidding schedule of producer i is

$$\beta_i(x) = \begin{cases} b_i(\underline{x}_i) & \text{for } x < \underline{x}_i \\ b_i(x) & \text{for } x \geq \underline{x}_i, \end{cases} \tag{6}$$

where $b_i(x)$ denotes the solution to (4) and $\underline{x}_i$ is the quantity they supply at the minimal demand level under this solution, i.e., $\underline{x}_i = D\big(b_i(\underline{x}_i), \underline{\varepsilon}\big) - (n-1)\beta_j^{-1}\big(b_i(\underline{x}_i)\big)$.

Bidding Behaviour in the Symmetric Equilibrium.— The symmetric equilibrium follows from using $x_i = x_j = x$, $p = \beta(x) = \beta_i(x) = \beta_j(x)$ and, consequently, $\beta^{-1'}(p) = 1/\beta'(x)$ in equations (4)–(6). Proposition (1) summarises the results.

Proposition 1 (Equilibrium under Pay-as-Bid Pricing) *Consider a pay-as-bid auction with $n \geq 2$ symmetric bidders with commonly known, weakly increasing marginal costs $c(x)$. Assume that market demand is given by some, weakly downward-sloping demand function $D(p, \varepsilon)$ such that $D_\varepsilon(\cdot) > 0$ where ε follows the cumulative distribution $F(\cdot)$ on $[\underline{\varepsilon}, \bar{\varepsilon}]$. The symmetric equilibrium in pure strategies is given by*

$$\beta^*_{PAB}(x) = \begin{cases} b^*_{PAB}(\underline{x}) & \textit{for } x < \underline{x} \\ b^*_{PAB}(x) & \textit{for } x \geq \underline{x}, \end{cases} \tag{7a}$$

with

$$b^{*'}_{PAB}(x) = \frac{(n-1)\, e_Q\big(b^*_{PAB}(x), nx\big)\, [b^*_{PAB}(x) - c(x)]\, f\Big(e\big(b^*_{PAB}(x), nx\big)\Big)}{1 - F\Big(e\big(b^*_{PAB}(x), nx\big)\Big) - e_p\big(b^*_{PAB}(x), nx\big)\, [b^*_{PAB}(x) - c(x)]\, f(\cdot)}, \tag{7b}$$

$$b^*_{PAB}(\bar{x}) = c(\bar{x}), \tag{7c}$$

$$\bar{x} = \frac{D(c(\bar{x}), \bar{\varepsilon})}{n}, \tag{7d}$$

and

$$\underline{x} = \frac{D(b^*_{PAB}(\underline{x}), \underline{\varepsilon})}{n}, \tag{7e}$$

where $e(p,Q)$ is such that $D\big(p,e(p,Q)\big)=Q$.

If demand is completely inelastic, we have that $e(p,Q)=Q$, in which case (7b) corresponds to the solution in Pycia and Woodward (2025).

Equilibrium in the Linear Model.— For the linear model from (2), we get $e(p,Q)=Q+m\,p$. In this case, there is a closed form for the symmetric equilibrium, which Corollary 1 presents.

Corollary 1 (Pay-as-Bid Pricing: Equilibrium in the Linear Model) *Consider the case where demand is $D(p,\varepsilon)=\varepsilon-mp$ with $\varepsilon\sim U[\underline{\varepsilon};\bar{\varepsilon}]$ and there are $n\geq 2$ symmetric producers, each with the commonly known marginal cost curve $c(x)=\alpha x$ with $\alpha,m>0$. The symmetric equilibrium in pure strategies in the pay-as-bid auction is given by the bidding function in* (7a) *and* (7e) *with*

$$b^*_{PAB}(x)=\underbrace{\frac{\sqrt{2\alpha m\,(2n-3)+(1-2n)^2+\alpha^2m^2}-2n+\alpha m+1}{4m}}_{\equiv\delta_{PAB}}\,x \\ -\underbrace{\frac{\bar{\varepsilon}\left(\sqrt{2\alpha m\,(2n-3)+(1-2n)^2+\alpha^2m^2}-2n-3\alpha m+1\right)}{4m\,(n+\alpha m)}}_{\equiv\lambda=\frac{(\alpha-\delta_{PAB})}{(n+m\alpha)}\,\bar{\varepsilon}}. \tag{8}$$

Bids under pay-as-bid pricing are less steep than marginal costs: Producers make high margins on low quantities, and their bids converge to their costs for high quantities. Mathematically, we can bound the slope of the increasing part of producers' bid schedules, $b^{*\prime}_{PAB}(x)=\delta_{PAB}$, by $(\alpha m-1)/2m<\delta_{PAB}<\alpha/2$.[5]

[5] To obtain these bounds, replace the term in the square root with $(\alpha m+2n-3)^2$ and $(\alpha m+2n-1)^2$, respectively.

3.3 Bidding Behaviour under Uniform Pricing

This section briefly revisits equilibrium behaviour in a uniform-price auction, which has been established by Klemperer and Meyer (1989). They show that – for given bids of the other suppliers – a bidder's equilibrium bids maximise the bidder's profits for every realisation of the demand-shifter ε, i.e., they are an ex-post equilibrium.

Bidder's Problem.— Therefore, we can model an individual producer i as maximising their profits for a given realisation of ε. If the other producers $j \neq i$ bid according to $\beta_j(p)$, producer i will supply $x_i = D(p, \varepsilon) - (n-1)\beta_j^{-1}(p)$ units at price p. Thus, they solve

$$\max_p \qquad \pi_{UP} = \big(D(p,\varepsilon) - (n-1)\beta_j^{-1}(p)\big)\, p - C\big(D(p,\varepsilon) - (n-1)\beta_j^{-1}(p)\big). \tag{9}$$

The first-order-condition of this problem is

$$x_i = \big(p - c(x_i)\big)\,\big((n-1)\beta_j^{-1\prime}(p) - D_p(p,\varepsilon)\big), \tag{10}$$

where $D_p(\cdot)$ is the derivative of $D(\cdot)$ with respect to p. The left-hand side of (10) reflects the upside of marginally raising the price, or equivalently, raising the bid on the x_i-th unit; the producer obtains additional marginal earnings on the x_i inframarginal units. However, they also lose marginal units and the corresponding margins, which the right-hand side of (10) reflects. Hence, producers bid truthfully on the first unit, as they do not benefit from increased margins on any inframarginal units. Proposition 2 describes bids in a symmetric equilibrium under uniform pricing based on the first-order condition in (10).

Proposition 2 (Klemperer and Meyer, 1989: Equilibrium under Uniform Pricing) *The bidding strategies in a symmetric equilibrium of a uniform-price auction with* $n \geq 2$ *symmetric bidders with commonly known, weakly increasing marginal costs* $c(x)$*, when market* *demand is given by some demand function $D(p, \varepsilon)$ satisfy the initial value problem*

$$\beta^{*\prime}_{UP}(x) = \frac{(n-1)[\beta^{*}_{UP}(x) - c(x)]}{x + D_p\big(\beta^{*}_{UP}(x)\varepsilon\big)[\beta^{*}_{UP}(x) - c(x)]} \tag{11a}$$

$$\beta^{*}_{UP}(0) = c(0). \tag{11b}$$

Proof. Use the definition of a symmetric equilibrium, i.e., $x_i = x_j = x$, $p = \beta(x) = \beta_i(x) = \beta_j(x)$, and $\beta^{-1\prime}(p) = 1/\beta'(x)$ in the first-order condition from (10). ■

Multiplicity of Equilibria and Equilibrium Selection.— Klemperer and Meyer (1989) show that there are infinitely many solutions to the initial value problem in (11), i.e., there are infinitely many symmetric equilibria under uniform pricing. Intuitively, producers' pay-offs only depend on a single point on their bid function for a given demand realisation; the rest of the bidding function stabilises the outcome (see Back and Zender, 2001).

When comparing price rules, this leaves a problem of equilibrium selection. Klemperer and Meyer (1989) show that there is (potentially) only one equilibrium, in which producers – for sufficiently high demand – (i) offer their entire, infinite production capacity and (ii) do so at weakly positive margins. The rest of this paper focuses on such equilibria. All other equilibria seem – at least to some extent – implausible in electricity markets. If (i) is violated, margins become negative at some point, resulting in negative profits for high demand levels.[6] If an equilibrium violates (ii), there is a strict upper bound to the number of units the capacity-unconstrained producers supply; for high enough demand levels, producers would even curb the quantity they supply as demand increases.[7] Competition authorities most likely would prevent such behaviour. Pycia and Woodward (2025) interpret the remaining equilbria as uncertainty-robust equilibria; no matter, how much demand potentially increases through the demand-shifter ε, i.e., independently of the distribution

---

[6]Of course, real-world producers do not have infinite production capacity. Thus, some of the equilibria that violate (i) might be plausible in the real world.

[7]Formally, equilibria violating (i) and (ii) correspond to the numerator and the denominator in (11a) going to zero such that the slope of bids goes to zero or infinity, respectively.

$F(\varepsilon)$, these equilibria remains economically meaningful. This robustness makes these equilibria focal for the comparison to pay-as-bid pricing.

Uncertainty-Robust Equilibrium in the Linear Model.— Corollary 2 presents the unique uncertainty-robust equilibrium for the linear model from (2).

Corollary 2 (Uniform Pricing: Uncertainty-Robust Equilibrium in the Linear Model) *Consider the case where demand is $D(p,\varepsilon) = \varepsilon - mp$ with $\varepsilon > 0$ and there are $n \geq 2$ symmetric producers, each with the commonly known marginal cost curve $c(x) = \alpha x$ with $\alpha, m > 0$. The only symmetric equilibrium in pure strategies in the uniform-price auction, in which bids are increasing for all $x \in [0,\infty)$ is given by the bidding function*

$$\beta^*_{UP}(x) = \underbrace{\frac{1}{2m}\left(2 + \alpha\, m - n + \sqrt{4 + \alpha^2 m^2 - 4n + 2\alpha mn + n^2}\right)}_{\equiv \delta_{UP}}\, x. \tag{12}$$

Corollary 2 is the generalisation of the linear example in Klemperer and Meyer (1989) to n producers. The proof is analogous to their proof for two suppliers.

Under uniform pricing, bids are steeper than marginal costs. By bounding the term in the square root by $(\alpha m + n - 2)^2$ and $(\alpha m + n)^2$, respectively, we get that the slope of bids under uniform pricing, δ_{UP}, satisfies $\alpha < \delta_{UP} < \alpha + 1/m$; for $n \to \infty$, $\delta_{UP} \to \alpha$.

The fact that bids under pay-as-bid pricing are less steep than marginal costs, and, therefore, than bids under uniform pricing immediately implies that price variance in terms of the stop-out price is lower under pay-as-bid pricing, as argued in previous literature (see Rassenti et al., 2003; Xiong et al., 2004). Furthermore, extreme price peaks, as mentioned in Section 1, are more likely under uniform pricing. In contrast to pay-as-bid pricing, the margins on the expensive-to-produce units that producers supply in high-demand situations are particularly high under uniform pricing.

4 Comparison of Market Outcomes

4.1 General Model: Perfect Competition and Constant Unit Costs

Pay-as-bid and uniform pricing result in first-best welfare and zero profits if marginal costs are constant or there is perfect competition in the sense that $n \to \infty$.

Perfect Competition.— At first sight, it is suprising that pay-as-bid pricing results in first-best welfare given that the slope of bids does not converge to that of marginal costs. However, with sufficiently many competitors, the quantity supplied by a single supplier converges to zero. Thus, we have $\bar{x} \to 0$ such that producers bid their true costs on the marginal unit they supply. Consequently, the two pricing rules are equivalent under perfect competition, as stated in Corollary 3.

Corollary 3 (Welfare & Consumer Surplus: Perfect Competition) *Consider $n \geq 2$ symmetric producers with the commonly known marginal cost curve $c(x)$, where $c(x) \geq 0$ and $c'(x) \geq 0$. The producers compete in a uniform-price or a pay-as-bid auction with market demand at price p given by the downward-sloping demand function $D(p, \varepsilon)$, where ε is a random variable. Under perfect competition, i.e., for $n \to \infty$ both auction formats implement the first-best welfare. As producers do not make profits, consumer surplus corresponds to first-best welfare.*

Proof. Under both pricing rules, the quantity supplied by an individual bidder, x_i, converges to zero as $n \to \infty$ for every demand realisation. It follows from Proposition 2 that the entire traded quantity is provided to consumers at price $\beta^*_{UP}(0) = c(0)$ in the uniform-price auction. Similarly, it follows from Proposition 1 that the entire traded quantity is provided at costs $\beta^*_{PAB}(\bar{x}) = c(\bar{x}) = c(0)$ to consumers under pay-as-bid pricing. The results on welfare and consumer surplus follow immediately. ■

The result of Corollary 3 contrasts with the findings of Federico and Rahman (2003), who conclude that the pay-as-bid auction does not implement first-best welfare under perfect competition. The difference to this paper is that their model softens competition by assuming that every producer corresponds to a point on the increasing industry cost curve. Thus, all but the marginal producer compete with competitors with higher costs. By contrast, they compete with equally competitive producers in this paper.

Constant Unit Costs.— With constant unit costs, producers bid their true costs under both pricing rules. For pay-as-bid pricing, the solution to (7b) with constant marginal costs is $b^*_{PAB}(x) = c(x)$, which implies that bidders truthfully reveal their costs for all units. Producers can only make positive margins to the degree that their competitors cannot undercut this price because they have higher costs when they produce more units. With constant marginal costs undercutting is always profitable and the pay-as-bid auction collapses into Bertrand competition. Similarly, it is easy to verify that truthful bidding is also an equilibrium under uniform pricing; it solves the initial value problem in (11). More precisely, truthful bidding is the uncertainty-robust equilibrium.

As stated in Corollary 4, these insights into equilibrium bidding behaviour directly imply that pay-as-bid and uniform pricing (in the uncertainty-robust equilibrium) implement the first-best welfare and all of this welfare goes to consumers.

Corollary 4 (Welfare & Consumer Surplus: Constant Marginal Costs) *Consider $n \geq 2$ symmetric producers with the commonly known marginal cost curve $c(x)$ with $c'(x) = 0$. The producers compete in a uniform-price or a pay-as-bid auction with market demand at price p given by the downward-sloping demand function $D(p, \varepsilon)$, where ε is a random variable. The pay-as-bid auction implements the first-best welfare. As producers have zero profits, consumer surplus corresponds to the first-best welfare. Uniform pricing results in the same market outcome in the uncertainty-robust equilibrium.*

Corollary 4 is the equivalent of the result in Genc (2008) for elastic demand.

4.2 Linear Model

The linear model allows deriving more detailed results on consumer surplus and welfare outside the extremely competitive cases from Section 4.1.

Consumer Surplus.— Figure 2 illustrates consumer surplus for different demand levels under both pricing rules. It shows that the comparison of consumer surplus becomes more favourable to pay-as-bid pricing as demand increases: First, bids are steeper under uniform pricing such that the stop-out price under uniform pricing increases faster in demand. Second, the rising stop-out price is more detrimental to consumer surplus under uniform pricing. In contrast to pay-as-bid pricing, consumers pay this increasing stop-out price on all units. As a result, consumer surplus is higher under pay-as-bid pricing for the highest demand realisation; under pas-as-bid pricing, bids are truthful on the last unit which implies a lower stop-out price under pay-as-bid pricing. Additionally, consumers even pay less than the stop-out price on inframarginal units. By contrast, uniform pricing can (but does not necessarily) result in higher consumer surplus when demand is low, as illustrated in Figure 2.

Proposition 3 claims that, in the expectation of all demand realisations, consumer surplus is higher under pay-as-bid than under uniform pricing. As is intuitive from Figure 2, the proof shows that the reason is that consumers pay less than the stop-out price on inframarginal units. Formally, consumer surplus under uniform pricing is $CS_{UP}(\varepsilon) = \int_0^{nx^*_{UP}} D^{-1}(x,\varepsilon)\,dx - nx^*_{UP}\,\beta^*_{UP}(x^*_{UP})$, where $D^{-1}(x,\varepsilon)$ is the inverse demand function for a given demand level ε and x^*_{UP} denotes the quantity supplied by an individual producer in equilibrium. For pay-as-bid pricing, consumer surplus is $CS_{PAB}(\varepsilon) = \int_0^{nx^*_{PAB}} D^{-1}(x,\varepsilon)\,dx - n\int_0^{x^*_{PAB}} \beta^*_{PAB}(x)\,dx$, where x^*_{PAB} is the equilibrium quantity supplied by an individual producer. According to Proposition 3, $\mathbb{E}[CS_{PAB}(\varepsilon)] > \mathbb{E}[CS_{UP}(\varepsilon)]$ in the linear model.

Proposition 3 (Linear Model: Higher Expected Consumer Surplus under Pay-as-Bid-Pricing) *Consider the case where demand is $D(p,\varepsilon) = \varepsilon - mp$ with $\varepsilon \sim U[\underline{\varepsilon};\bar{\varepsilon}]$ and there are $n \geq 2$ symmetric producers, each with the commonly known marginal cost curve*

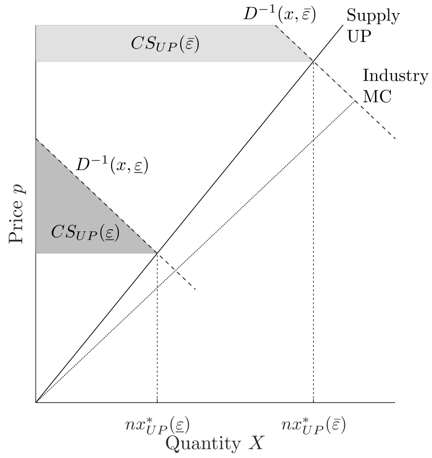

(a) Uniform pricing.

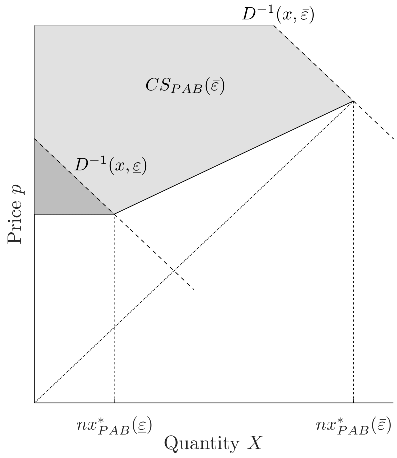

(b) Pay-as-bid pricing.

Figure 2: Consumer surplus under uniform and pay-as-bid pricing. The dotted and solid lines represent industry-wide marginal costs and supply, respectively; the darkly and lightly shaded areas correspond to the consumer surplus for the lowest and highest demand realisation, respectively.

$c(x) = \alpha x$ with $\alpha, m > 0$. Without loss of generality, assume that $m = 1$. Expected consumer surplus is strictly higher under pay-as-bid pricing than in the uncertainty-robust equilibrium of the uniform-price auction assuming that the traded quantity in the pay-as-bid auction is positive, i.e., that $\underline{\varepsilon} \geq b^*_{PAB}(0)$.

Proof. See Appendix A. ■

The result of Proposition 3 is not surprising in the context of theoretical literature. For example, Pycia and Woodward (2025) show that for inelastic demand the *optimally designed* pay-as-bid auction results in a higher expected consumer surplus than the *optimally designed* uniform price auction without making any substantial functional form assumptions.[8] The big difference to the result in Proposition 3 is that they have the auctioneer choose the distribution of demand (and stochastic reserve prices) to maximise the expected consumer surplus under either pricing rule. Thus, their result does not say much about the policy choice

[8] This is, translating their results to the reverse auction.

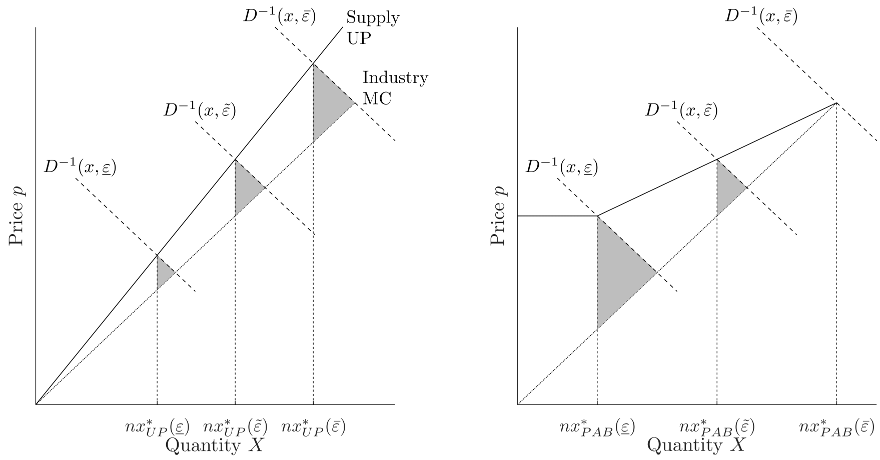

(a) Uniform pricing. (b) Pay-as-bid pricing.

Figure 3: Welfare under uniform and pay-as-bid pricing in the linear model. The dotted and solid lines represent industry-wide marginal costs and supply, respectively; the shaded triangles correspond to the deadweight loss for different demand levels.

between pricing rules in electricity markets: Regulators in the context of electricity markets have to take demand and its variation as given. Proposition 3 speaks to that situation.

Welfare.— The comparison of expected welfare is generally ambiguous. Comparing welfare between the two pricing rules is equivalent to comparing the deadweight losses, i.e., the welfare loss compared to the first-best. Figure 3 depicts the deadweight loss for varying demand levels under both pricing rules as shaded triangles.

As illustrated in Figure 3, the deadweight loss under pay-as-bid pricing becomes smaller as demand increases, whereas it increases with demand under uniform pricing. Pay-as-bid pricing eliminates the deadweight loss when demand is at its maximum because the stop-out price corresponds to marginal costs. Consequently, pay-as-bid pricing results in higher welfare than uniform pricing if demand is close to its highest level. For lower demand levels, the producers' margins increase under pay-as-bid pricing and so does the deadweight loss. By contrast, bids are steeper than marginal costs for uniform pricing such that margins increase with demand, resulting in an increasing deadweight loss.

Figure 3 makes clear that the deadweight loss under pay-as-bid and uniform pricing is $DWL_{PAB/UP} = 0.5 \cdot (\beta^*_{PAB/UP}(x^*_{PAB/UP}) - \alpha\, x^*_{PAB/UP}) \cdot (n\, x^*_{PAB/UP} - n\, x_{FB})$, where $x_{FB} = \varepsilon/(n + m\alpha)$ is the quantity an individual producer supplies in the first-best. When $x^*_{PAB} = x^*_{UP}$ both pricing rules result in the same welfare, which is the case for the demand level

$$\tilde{\varepsilon} = \underbrace{\frac{(\alpha - \delta_{PAB})(n + \delta_{UP})}{(\delta_{UP} - \delta_{PAB})(n + \alpha)}}_{\equiv \tilde{\mu} < 1} \bar{\varepsilon}. \tag{13}$$

Pay-as-bid pricing results in higher welfare if and only if $\varepsilon > \tilde{\varepsilon}$ and uniform pricing outperforms when $\varepsilon < \tilde{\varepsilon}$. Equation (13) shows that $\tilde{\varepsilon}$ is always strictly lower than $\bar{\varepsilon}$, which confirms that there are always some (high-demand) situations in which pay-as-bid pricing does better than uniform pricing in terms of welfare. Interestingly, none of the terms on the right-hand side of equation (13) depend on $\underline{\varepsilon}$. Thus, there is no guarantee that there is a demand level for which uniform pricing results in higher welfare than pay-as-bid pricing; we can have $\underline{\varepsilon} > \tilde{\varepsilon}$, in which case pay-as-bid pricing results in higher welfare for every single demand realisation. This is the case for any demand distribution for which the ratio of the minimum and maximum demand levels, $\mu = \underline{\varepsilon}/\bar{\varepsilon}$, is sufficiently large; i.e., that has $\mu > \tilde{\mu}$, where – according to equation (13) – $\tilde{\mu}$ does not depend on the demand distribution.

Proposition 4 shows that something similar holds for expected welfare. Pay-as-bid pricing results in higher expected welfare than uniform pricing if and only if $\mu > \tilde{\mu}^E$, where $\tilde{\mu}^E$ does not depend on the demand distribution: An increase in μ is an increase of the lowest demand level $\underline{\varepsilon}$ for a given maximum demand level $\bar{\varepsilon}$; thus, it relatively raises the expected welfare under pay-as-bid pricing by shifting probability mass towards those cases where pay-as-bid pricing does better. As higher *expected* welfare under pay-as-bid pricing does not require that pay-as-bid pricing results in higher welfare for every demand realisation, $\tilde{\mu}^E < \tilde{\mu}$.

The upshot is that changes in the demand distribution affect the welfare comparison if and only if they change μ. Intuitively, μ captures demand uncertainty by measuring the distance between the minimum and maximum demand levels multiplicatively. Formally, μ

fully determines and is negatively related to the geometric variance of the demand level ε, or, equivalently, the variance of $\log(\varepsilon)$. Thus, μ is a very specific measure of demand uncertainty; it measures the extent to which ε multiplicatively deviates from its mean.[9] Proposition 4 formalises the results regarding the comparison of (expected) welfare in the linear model.

Proposition 4 (Linear Model: Expected Welfare) *Consider the case where demand is* $D(p, \varepsilon) = \varepsilon - mp$ *with* $\varepsilon \sim U[\underline{\varepsilon}; \bar{\varepsilon}]$, *where* $\underline{\varepsilon} = \mu\bar{\varepsilon}$ *with* $0 < \mu < 1$. *There are* $n \geq 2$ *symmetric producers, each with the commonly known marginal cost curve* $c(x) = \alpha x$ *with* $\alpha, m > 0$, *that compete either in a pay-as-bid or a uniform-price auction. Without loss of generality, assume that* $m = 1$. *Furthermore, assume that the traded quantity in the pay-as-bid auction is positive, i.e., that* $\mu > \underline{\mu} = \frac{(\alpha - \delta_{PAB})}{(\alpha + n)}$. *If* $\mu > \tilde{\mu}$, *welfare is higher under pay-as-bid pricing for every single demand realisation, i.e.,*

$$\mu > \frac{(\alpha - \delta_{PAB})(n + \delta_{UP})}{(\delta_{UP} - \delta_{PAB})(n + \alpha)} \qquad \Rightarrow \qquad W^*_{PAB}(\varepsilon) > W^*_{UP}(\varepsilon) \quad \forall \varepsilon \in [\underline{\varepsilon}; \bar{\varepsilon}].$$

By contrast, a complete welfare-dominance of uniform pricing is impossible. The critical demand level for the dominance of the pay-as-bid auction, $\tilde{\mu}$, *is increasing in* α. *For a given maximum level of demand* $\bar{\varepsilon}$, *the expected welfare under pay-as-bid pricing rises relative to uniform pricing when the minimum demand level increases, i.e.,* $\frac{\partial \mathbb{E}[\Delta W]}{\partial \mu} > 0$, *where* $\mathbb{E}[\Delta W] = \mathbb{E}[W^*_{PAB}] - \mathbb{E}[W^*_{UP}]$. *Thus, if the minimal demand level is sufficiently high, pay-as-bid pricing results in higher expected welfare, i.e.,*

$$\mu > \max\left\{\underline{\mu}; \tilde{\mu}^E\right\} \qquad \Rightarrow \qquad \mathbb{E}\left[W^*_{PAB}\right] > \mathbb{E}\left[W^*_{UP}\right],$$

[9]By contrast, the arithmetic variance of ε – which measures the additive deviation of ε from its mean – depends on μ *and* $\bar{\varepsilon}$. Hence, depending on $\bar{\varepsilon}$, the two pricing rules can perform differently regarding (expected) welfare for two distributions with the same arithmetic variance.

where $\mu > \underline{\mu}$ corresponds to the assumption of positive trade under pay-as-bid pricing and

$$\tilde{\mu}^E = \frac{-\left(\frac{(\delta_{UP}-\alpha)^2}{(n+\delta_{UP})^2} + \frac{2(\alpha-\delta_{PAB})^2}{(n+\delta_{PAB})^2}\right) + \frac{(\delta_{UP}-\alpha)}{(n+\delta_{UP})}\sqrt{\frac{12(\alpha-\delta_{PAB})^2}{(n+\delta_{PAB})^2} - \frac{3(\delta_{UP}-\alpha)^2}{(n+\delta_{UP})^2}}}{2\left(\frac{(\delta_{UP}-\alpha)^2}{(n+\delta_{UP})^2} - \frac{(\alpha-\delta_{PAB})^2}{(n+\delta_{PAB})^2}\right)} < \tilde{\mu}. \tag{14}$$

If $n \leq 3$ and $\alpha < (4-n)/4 < m$, we have $\underline{\mu} > \tilde{\mu}^E$, such that pay-as-bid pricing results in higher expected welfare for any uniform distribution with positive trade, i.e. for all $\mu \in [\underline{\mu}; 1)$.

Proof. See Appendix B. ■

In addition to the previously discussed, Proposition 4 claims that a complete welfare-dominance of pay-as-bid pricing is "more likely," i.e., can be realised for more parametrisations of the demand distribution if marginal costs are relatively flat. Similarly, it shows that pay-as-bid pricing results in higher expected welfare for *any* (uniform) distribution of demand if there are few suppliers with relatively flat marginal costs. Regarding the slope of marginal costs, Federico and Rahman (2003) obtain a similar result in a monopoly.

5 Implications & Discussion

Even though the model of this paper necessarily abstracts from the full complexity of real-world electricity markets, it nonetheless has important policy implications.

Consumer Surplus.— First and foremost, this paper suggests that consumers are better off under pay-as-bid pricing. Thus, to the extent public policy attaches particular value to raising consumer surplus, it should consider moving from uniform to pay-as-bid pricing. Such a policy focus seems worthwhile given limited competition in electricity markets.

Welfare.— However, this paper also confirms that this increase in consumer surplus is not necessarily driven by an overall increase in welfare. Hence, a switch to pay-as-bid pricing is not an obvious decision. In the model, lower demand uncertainty favours pay-as-bid pricing in terms of welfare. Consequently, experiments with pay-as-bid pricing should start on markets where most of the uncertainty is already resolved. For example, pay-as-bid pricing is more promising in short-term markets (e.g., the spot market) than in longer-term markets.

Strategic Tractability.— Proposition 1 shows that optimal behaviour under pay-as-bid pricing is tractable under fairly general conditions. Hence, switching to pay-as-bid pricing is unlikely to result in erratic behaviour of market participants. If anything, such erratic behaviour seems more likely under uniform pricing, which allows for multiple equilibria. Furthermore, bringing data to the model allows realistic predictions on welfare under either pricing rule in a way that is firmly grounded in economic theory. This way, future research could substantially lower policymakers' uncertainty regarding the effects of price rules.

Price Peaks and Price Variance.— Last, the fact that bid functions are steeper under uniform pricing implies that price variance, in terms of the stop-out price, is higher under uniform pricing. In particular, uniform pricing exacerbates price peaks because producers make particularly high margins when using high-cost technologies. To the degree that price variance is an economy-wide bad as it complicates planning, this is a further argument for pay-as-bid pricing.

6 Conclusion

Previous theoretical literature on the comparison of uniform and pay-as-bid pricing in electricity markets has made strong simplifying assumptions, ranging from constant unit costs over completely inelastic demand to abstracting from the oligopolistic nature of the market. This paper has derived the equilibrium strategies in the pay-as-bid auction for an oligopoly of symmetric suppliers with an arbitrary, increasing marginal cost function and an arbitrary random, downward-sloping, demand. This general model showed that both pricing rules implement the first-best welfare with zero profits for suppliers when competition is intense because marginal costs are constant or there are infinitely many suppliers.

To derive more detailed results, this paper has also considered a linear model, where marginal costs and demand are linear, and the intercept of the demand curve follows a uniform distribution. In the linear model, pay-as-bid pricing results in a higher expected consumer surplus because it lowers prices on inframarginal units. By contrast, the comparison of expected welfare between the unique equilibrium under pay-as-bid pricing and the uncertainty-robust equilibrium of the uniform-price auction, i.e., the only equilibrium that remains meaningful for any distribution of demand, is ambiguous. However, pay-as-bid pricing outperforms uniform pricing for at least some demand realisations. It results in higher expected welfare if demand variability is sufficiently low and can even outperform uniform pricing for all demand realisations if demand variability is very low.

This paper shows that optimal bidding is tractable – and, thus, in principle, predictable – under pay-as-bid pricing under general conditions; the risk of strategic bidding should not keep regulators from considering it. A switch to pay-as-bid pricing presumably would benefit consumers, whereas its effects on welfare are unclear. Positive welfare effects are most likely in short-term markets with low demand uncertainty.

A Proof of Proposition 3

The equilibrium under uniform pricing maximises consumer surplus conditional on the pricing rule and bids: A unit is traded if and only if the value it creates for consumers is larger than the associated price. In contrast to uniform pricing, trading a unit does not impact the price of other units such that the equilibrium quantity maximises consumer surplus. Therefore, to prove Proposition 3, it is sufficient to show that $CS_{PAB}(nx^*_{UP}) > CS^*_{UP}$ given that $CS^*_{PAB} > CS_{PAB}(nx^*_{UP})$. As both pricing rules would generate the same gross value to consumers when trading nx^*_{UP}, we can prove Proposition 3 by showing that consumers' total payments would be lower under pay-as-bid pricing when trading nx^*_{UP}, i.e., that

$$\Delta P = \underbrace{\mathbb{E}\Big[nx^*_{UP}\ \beta^*_{UP}\big(x^*_{UP}(\varepsilon)\big)\Big]}_{\text{Expected Payment in UP}} - \underbrace{\mathbb{E}\Big[n\int_0^{x^*_{UP}(\varepsilon)} \beta^*_{PAB}(x)\ dx\Big]}_{\text{Exp. Pay. in PAB when trading } nx^*_{UP}} > 0. \tag{A1}$$

Because we do not know, for what demand levels $x^*_{UP}(\varepsilon) > x^*_{PAB}(\underline{\varepsilon})$ such that the increasing part of the bidding function $\beta^*_{PAB}(\cdot)$ becomes relevant. We can overestimate payments under pay-as-bid pricing by assuming that $x^*_{UP}(\varepsilon) > x^*_{PAB}(\underline{\varepsilon})$, i.e.,

$$\mathbb{E}\Big[n\int_0^{x^*_{UP}(\varepsilon)} \beta^*_{PAB}(x)\ dx\Big] < \mathbb{E}\left[n\ \left(\int_0^{x^*_{UP}(\varepsilon)}\Big(\lambda + \delta_{PAB}\,x\Big)dx + 0.5\,\delta_{PAB}\ \big(x^*_{PAB}(\underline{\varepsilon})\big)^2\right)\right],$$

where $\lambda = b^*_{PAB}(0) = (\alpha - \delta_{PAB})x^*_{PAB}(\bar{\varepsilon}) = \frac{(\alpha-\delta_{PAB0})}{(n+m\alpha)}\,\bar{\varepsilon}$. Using this bound in (A1) and dividing by n, we get that for $\Delta P > 0$ it is sufficient that

$$\frac{(2\delta_{UP} - \delta_{PAB})}{2(n+\delta_{UP})^2}\ \mathbb{E}[\varepsilon^2] - \frac{2\lambda}{2(n+\delta_{UP})}\ \mathbb{E}[\varepsilon] - \frac{\delta_{PAB}\,(\underline{\varepsilon}-\lambda)^2}{2(n+\delta_{PAB})^2} > 0, \tag{A2}$$

where we used $x^*_{UP}(\varepsilon) = \frac{\varepsilon}{n+m\delta_{UP}}$. Rewriting the expected values by setting $\underline{\varepsilon} = \mu\bar{\varepsilon}$, where $\mu \in [\underline{\mu}; 1)$ with $\underline{\mu} = \lambda/\bar{\varepsilon} = \frac{(\alpha-\delta_{PAB})}{(\alpha+n)}$ by the assumption of positive trade, i.e., $\underline{\varepsilon} > \lambda$ and dividing by $(\bar{\varepsilon})^2$, we get

$$\Delta\tilde{P} = \frac{(2\delta_{UP} - \delta_{PAB})\,(1+\mu+\mu^2)}{6(n+\delta_{UP})^2} - \frac{(\alpha-\delta_{PAB})\,(1+\mu)}{2(n+\delta_{UP})\ (n+\alpha)} - \frac{\delta_{PAB}}{2(n+\delta_{PAB})^2}\left(\mu - \frac{(\alpha-\delta_{PAB})}{(\alpha+n)}\right)^2 > 0, \tag{A3}$$

which is equivalent to (A2) under the assumptions of Proposition 3. Given that (A3) is a polynomial of degree to 2 in μ, (A3) holds if (a) $\Delta\tilde{P}(\underline{\mu}) > 0$, (b) $\Delta\tilde{P}'(\underline{\mu}) > 0$ and (c) $\Delta\tilde{P}(1) > 0$.

Proof of (a): For $\Delta\tilde{P}(\underline{\mu})$, we get

$$\frac{(2\delta_{UP} - \delta_{PAB})\,(1+\underline{\mu}+\underline{\mu}^2)}{6(n+\delta_{UP})^2} - \frac{(\alpha-\delta_{PAB})\,(1+\underline{\mu})}{2(n+\delta_{UP})\ (n+\alpha)} > 0$$

$$\overset{\substack{\delta_{PAB}<\alpha/2\\ \delta_{PAB}>(\alpha-1)/2}}{\Leftarrow}\ \frac{(2\delta_{UP} - \delta_{PAB})}{6(n+\delta_{UP})^2} - \frac{(\alpha-\delta_{PAB})}{2(n+\delta_{UP})\ (\alpha+n)} + \frac{(2\delta_{UP} - \delta_{PAB})}{6(n+\delta_{UP})^2}\ \frac{\alpha^2}{2(\alpha+n)}\ \frac{1}{(3\alpha+2n+1)} > 0$$

$$\overset{\substack{n>1 \\ \delta_{PAB}<\alpha/2}}{\Leftarrow} \quad (2\delta_{UP} - \alpha/2)\,\frac{6(\alpha+n)^2+\alpha^2}{6(\alpha+n)} - 3(n+\delta_{UP})(\alpha-\delta_{PAB}) > 0$$

$$\overset{\substack{\delta_{UP}>\alpha \\ \delta_{PAB}\delta_{UP}\geq\frac{\alpha^2}{2}}}{\Leftarrow} \quad (2n-\alpha)\delta_{UP} + 3n\delta_{PAB} + a^2 - \frac{7n}{2}\alpha + \frac{\alpha^3}{4(\alpha+n)} > 0, \tag{A4}$$

To prove that (A4) holds, we transform it into easy-to verify polynomials: For $n = 2$, (A4) resolves to

$$\alpha^2 - 7\alpha + \frac{a^3}{2(a+2)} + (4-\alpha)\sqrt{\alpha(4+\alpha)} + 3\sqrt{9+\alpha(2+\alpha)} - 9 > 0.$$

For $n \geq 3$, use $\delta_{PAB} > \frac{(n-1)}{(2n-1)}\alpha$, to obtain the sufficient condition

$$(2n-\alpha)\delta_{UP} + \alpha^2 + \frac{\alpha^3}{4(a+n)} - \frac{(8n-1)n\alpha}{2(2n-1)} > 0. \tag{A5}$$

To reduce this problem to a simple polynomial, we can bound δ_{UP} distinguishing $2n \lesseqqgtr \alpha$. *If* $2n < \alpha$, δ_{UP} enters negatively, and we get a sufficient condition by using $\delta_{UP} < \alpha + 1$. Hence, a sufficient condition for (A5) in this case, is

$$\alpha^3 + \frac{(4-14n)}{(2n-1)}\,\alpha^2 + \frac{2n(n-2)}{(2n-1)}\,\alpha + 8n^2 > 0.$$

If $2n \geq \alpha$, we can use the additional restriction on α to give a particularly tight lower bound on δ_{UP}, i.e., $\delta_{UP} > \alpha + \frac{277(n+2)}{400n(n+\alpha)}\,\alpha$ to rewrite the sufficient condition from (A5) as

$$100n(2n-1)\alpha^2 + (-1154n^2 - 831n + 554)\alpha + 508n^3 + 1662n^2 - 1108n.$$

Proof of (b): $\Delta\tilde{P}'(\underline{\mu}) > 0$ follows directly from (a). We have,

$$\Delta\tilde{P}'(\underline{\mu}) = \frac{(2\delta_{UPA} - \delta_{PABA})\,(1+2\underline{\mu})}{6(n+\delta_{UPA})^2} - \frac{(\alpha - \delta_{PABA})}{2(n+\delta_{UPA})\,(n+\alpha)},$$

which is larger than $\Delta\tilde{P}(\underline{\mu})$ given that $\underline{\mu} > 0$ and $\underline{\mu}^2 < \underline{\mu}$ as $\underline{\mu} < 1$.
Proof of (c): We need to show that

$$\Delta\tilde{P}(1) = \frac{(2\delta_{UP} - \delta_{PAB})}{2(n+\delta_{UP})^2} - \frac{(\alpha - \delta_{PAB})}{(n+\delta_{UP})\,(n+\alpha)} - \frac{\delta_{PAB}}{2(n+a)^2} > 0.$$

$\Delta\tilde{P}(1)$ is decreasing in δ_{PAB}. Thus, $\Delta\tilde{P}(1)$ is strictly positive as long as $\delta_{PAB} < \frac{2n(\alpha+n)}{(\delta_{UP}-\alpha)}$, which is satisfied as $(\delta_{UP} - \alpha)\,\delta_{PAB} < (\alpha + 1 - \alpha)\,\delta_{PAB} = \delta_{PAB} < \alpha < 2n(\alpha+n)$.

B Proof of Proposition 4

Critical Value for Dominance of Pay-as-Bid Pricing $\tilde{\mu}$ decreases in α.— Proposition 4 claims that $\frac{\partial \tilde{\mu}}{\partial \alpha} > 0$, or equivalently,

$$-\left(\delta_{UP} - \delta_{PAB}\right) + \frac{(n+\alpha)}{(n+\delta_{UP})}\left(\alpha - \delta_{PAB}\right)\frac{\partial \delta_{UP}}{\partial \alpha} + \frac{(n+\alpha)}{(n+\delta_{PAB})}\left(\delta_{UP} - \alpha\right)\frac{\partial \delta_{PAB}}{\partial \alpha} < 0. \quad \text{(B1)}$$

A sufficient condition for (B1) to hold follows from replacing the first factors in the second and third summand by one, which, after rewriting the derivatives, gives

$$-(\delta_{UP}-\delta_{PAB})+(\alpha-\delta_{PAB})\frac{(\delta_{UP}+n-1)}{(2\delta_{UP}-a+n-2)}+(\delta_{UP}-\alpha)\frac{(\delta_{PAB}+n-1)}{(4\delta_{PAB}-a+2n-1)} < 0. \quad \text{(B2)}$$

For $\alpha < 1/2$, use that $\frac{(n-1)}{(2n-1)}\alpha < \delta_{PAB} < \frac{\alpha}{2} - \frac{\alpha}{2(2n-1+\alpha)}$ in the first two summands, and $a/4 < \delta_{PAB} < a/2$ in the third summand. After multiplying by $2\sqrt{\alpha^2 + 2\alpha n + n^2 - 4n + 4}$, adding $(1/4 - a/2)(a^2/(2n-1)) > 0$, and multiplying by $2(2n-1)(2n-1+\alpha)$, we get that a sufficient condition for (B2) is

$$\begin{aligned}\underbrace{\left(2n+\frac{1}{2}\right)}_{>0}\alpha^3 + \underbrace{\left(3n^2 - 5n + \frac{7}{2}\right)}_{>0}\alpha^2 + \underbrace{\left(-4n^3 + n^2 + 4n - 4\right)}_{<0}\alpha - 2n(n-2)^2(2n-1) \\ < \left(a^3 + (3n-5)\alpha^2 - (n-2)\alpha - n(4n^2 - 10n + 4)\right)\underbrace{\sqrt{\alpha^2 + 2\alpha n + n^2 - 4n + 4}}_{>0}.\end{aligned} \quad \text{(B3)}$$

The left-hand side is negative; the first two summands are dominated by the third one given that $\alpha < 1/2$. The right-hand side is negative if and only if $n \geq 3$. Thus, (B3) holds for $n = 2$. For $n \geq 3$, squaring both sides and dividing by $\alpha^2 > 0$, (B3) is equivalent to

$$\begin{aligned}0 > \alpha^6 + (8n-10)\alpha^5 + \underbrace{\frac{(72n^2 - 232n + 131)}{4}}_{>0}\alpha^4 + \frac{(8n^3 - 134n^2 + 254n - 127)}{2}\alpha^3 \\ + \frac{(-96n^4 + 360n^3 + 140n^2 + 447)}{4}\alpha^2 + (-16n^5 + 140n^4 - 264n^3 + 149n^2 + 52n - 52)\alpha \\ - (8n^5 - 176n^4 + 490n^3 - 408n^2 + 104n),\end{aligned}$$

where we can obtain an upper bound for the right-hand side by using the fact $\alpha < 1$, i.e., it is sufficient that

$$\begin{aligned}0 > \alpha^2 + (8n-10)\alpha^2 + \frac{(72n^2 - 232n + 131)}{4}\alpha^2 + \frac{(8n^3 - 134n^2 + 254n - 127)}{2}\alpha^3 \\ + \frac{(-96n^4 + 360n^3 + 140n^2 + 447)}{4}\alpha^2 + (-16n^5 + 140n^4 - 264n^3 + 149n^2 + 52n - 52)\alpha \\ - (8n^5 - 176n^4 + 490n^3 - 408n^2 + 104n)\alpha.\end{aligned}$$

Replacing α by $2\alpha^2\ \alpha$ in the last two summands, after dividing by α^2, it is sufficient to show that, for $\alpha < 1/2$ and $n \geq 3$,

$$0 > (8n^3 - 134n^2 + 254n - 127)\,\alpha + (-96n^5 - 192n^4 + 1084n^3 - 930n^2 + 122n + 63).$$

For $\alpha \geq 1/2$, a sufficient condition for (B2) is $-(\delta_{UP} - \delta_{PAB}) + (\alpha - \delta_{PAB})\ \frac{\partial \delta_{UP}}{\partial \alpha} + \frac{(\delta_{UP}-\alpha)}{2} < 0$, given that $\delta_{PAB} > (a-1)/2$, or equivalently,

$$(\alpha - \delta_{PAB}) - \left(\frac{(n-2) + 2(\delta_{UP} - \delta_{PAB}) + \alpha}{2}\right)(\delta_{UP} - \alpha) < 0.$$

Using the bounds $\frac{(n-1)}{(2n-1)}\alpha < \delta_{PAB} < \alpha/2$ for the first and second δ_{PAB}, respectively, we get the alternative sufficient condition

$$(n-2)^2 + \left(\frac{2n^2 - n - 2}{(2n-1)}\right)\alpha > (n-2)\sqrt{\alpha^2 + 2\alpha n + n^2 - 4n + 4}. \tag{B4}$$

For $n = 2$, (B4) evidently holds, as only the positive second summand on the left-hand side is non-zero. For $n \geq 3$ both the left- and right-hand side of (B4) are positive. Thus, after squaring both sides, (B4) is equivalent to

$$\left(\left(\frac{2n^2 - n - 2}{(2n-1)(n-2)}\right)^2 - 1\right)\alpha - \frac{4}{(2n-1)} > 0.$$

Relative Expected Welfare under Pay-as-Bid Pricing Increases with Higher Minimal Demand.— The difference in welfare between the pricing rules for a given demand realisation $\varepsilon = z\bar{\varepsilon}$ with $z \in [\underline{\mu}; 1)$, where from the assumption of positive trade $\underline{\mu} = \frac{(\alpha - \delta_{PAB})}{(\alpha+n)}$, is the difference in the respective deadweight loss, i.e.,

$$\Delta W = (z\bar{\varepsilon})^2 \frac{n}{2} \underbrace{\left(\frac{1}{(n+m\alpha)} - \frac{1}{(n+\delta_{UP})}\right)}_{>0} \underbrace{\left(\frac{(\delta_{UP} - \alpha)}{(n+\delta_{UP})}\right)}_{>0}$$
$$- (\bar{\varepsilon}^2)\frac{n}{2} \underbrace{\left(\frac{z}{(n+m\alpha)} - \frac{(z-\tilde{\lambda})}{(n+\delta_{PAB})}\right)}_{>0} \underbrace{\left(\frac{(\delta_{PAB} - \alpha)(z-\tilde{\lambda})}{(n+\delta_{PAB})} + \tilde{\lambda}\right)}_{>0},$$

where $\tilde{\lambda} = \lambda/\bar{\varepsilon} = \frac{(\alpha - \delta_{PAB})}{(\alpha+n)}$. The derivative of ΔW with respect to μ is

$$\frac{\partial \Delta W}{\partial z} = z(\bar{\varepsilon})^2\, n \underbrace{\left(\frac{1}{(n+m\alpha)} - \frac{1}{(n+\delta_{UP})}\right)}_{>0} \underbrace{\left(\frac{(\delta_{UP} - \alpha)}{(n+\delta_{UP})}\right)}_{>0} - (\bar{\varepsilon}^2)\frac{n}{2} \underbrace{\left(\frac{z}{(n+m\alpha)} - \frac{(z-\tilde{\lambda})}{(n+\delta_{PAB})}\right)}_{>0} \underbrace{\left(\frac{(\delta_{PAB} - \alpha)}{(n+\delta_{PAB})}\right)}_{<0}$$
$$- (\bar{\varepsilon}^2)\,\frac{n}{2} \underbrace{\left(\frac{1}{(n+m\alpha)} - \frac{1}{(n+\delta_{PAB})}\right)}_{<0} \underbrace{\left(\frac{(\delta_{PAB} - \alpha)(z-\tilde{\lambda})}{(n+\delta_{PAB})} + \tilde{\lambda}\right)}_{>0} > 0.$$

From $\frac{\partial \Delta W}{\partial z} > 0$, it follows that $\frac{\partial \mathbb{E}[\Delta W]}{\partial \mu} > 0$: Expected welfare can be written as

$$\mathbb{E}[\Delta W] = \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} \Delta W(\varepsilon; \bar{\varepsilon}) \frac{1}{\bar{\varepsilon} - \underline{\varepsilon}} \, d\varepsilon = \bar{\varepsilon}^2 \int_{\mu}^{1} \Delta \tilde{W}(z) \frac{1}{(1-\mu)} \, dz, \tag{B5}$$

where $\Delta \tilde{W}(z) = \Delta W / (\bar{\varepsilon})^2$. Hence, $\frac{\mathbb{E}[\Delta W]}{\partial \mu} = \frac{\bar{\varepsilon}^2}{(1-\mu)} \left(\frac{1}{(1-\mu)} \int_{\mu}^{1} \Delta \tilde{W}(z) \, dz - \Delta \tilde{W}(\mu) \right) > 0$.

Conditions for Higher Expected Welfare under Pay-as-Bid Pricing.— Based on (B5), pay-as-bid pricing results in higher expected welfare if and only if $\int_{\mu}^{1} \Delta \tilde{W}(z) \, dz > 0$. Given that $\frac{\partial \mathbb{E}[\Delta W]}{\partial \mu} > 0$ the critical value of μ for (B6) to hold is $\tilde{\mu}^E < \tilde{\mu}$. Formally $\tilde{\mu}^E$ solves

$$\int_{\tilde{\mu}^E}^{1} \Delta \tilde{W}(z) \, dz = 0, \tag{B6}$$

or equivalently, $\frac{(\delta_{UP}-\alpha)^2}{(n+\delta_{UP})^2} \left(1 - (\tilde{\mu}^E)^3\right) - \frac{(\alpha-\delta_{PAB})^2}{(n+\delta_{PAB})^2} (1 - \tilde{\mu}^E)^3 = 0$. Given that $\tilde{\mu}^E = 1$ is an obvious solution to (B6) but not the desired solution with $\tilde{\mu}^E < \tilde{\mu}$, after polynomial long division with $(1 - \tilde{\mu}^E)$, we have that

$$\left(\frac{(\delta_{UP} - \alpha)^2}{(n + \delta_{UP})^2} - \frac{(\alpha - \delta_{PAB})^2}{(n + \delta_{PAB})^2} \right) \left((\tilde{\mu}^E)^2 + 1 \right) + \left(\frac{(\delta_{UP} - \alpha)^2}{(n + \delta_{UP})^2} + \frac{2(\alpha - \delta_{PAB})^2}{(n + \delta_{PAB})^2} \right) (\tilde{\mu}^E) = 0. \tag{B7}$$

Distinguish two cases: If $\frac{(\delta_{UP}-\alpha)^2}{(n+\delta_{UP})^2} \geq \frac{(\alpha-\delta_{PAB})^2}{(n+\delta_{PAB})^2}$ the left-hand side of (B7) is positive for all $\tilde{\mu}^E > 0$, i.e., expected welfare under pay-as-bid pricing exceeds expected welfare under uniform pricing for any μ.[10] If $\frac{(\delta_{UP}-\alpha)^2}{(n+\delta_{UP})^2} < \frac{(\alpha-\delta_{PAB})^2}{(n+\delta_{PAB})^2}$ (B7) has one solution for which $\tilde{\mu}^E < 1$, which is given in (14).

Cases in which Expected Welfare is Higher under Pay-as-Bid Pricing for Any Level of Demand Uncertainty.— If $\frac{(\delta_{UP}-\alpha)^2}{(n+\delta_{UP})^2} > \frac{(\alpha-\delta_{PAB})^2}{(n+\delta_{PAB})^2}$ pay-as-bid pricing results in higher welfare for any $\mu \in [\underline{\mu}; 1)$. Given that $\delta_{PAB} > \frac{(n-1)\alpha}{(2n-1)}$ for this to be the case, it is sufficient that

$$\frac{(\delta_{UP} - \alpha)}{(n + \delta_{UP})} > \frac{\left(\alpha - \frac{(n-1)\alpha}{(2n-1)}\right)}{\left(n + \frac{(n-1)\alpha}{(2n-1)}\right)} = \frac{\alpha n}{(2n^2 + (\alpha - 1)n - \alpha n)} > \frac{\alpha}{(2n-1)},$$

if $\alpha < 1$. The sufficient condition $\frac{(\delta_{UP}-\alpha)}{(n+\delta_{UP})} > \frac{\alpha}{(2n-1)}$ is equivalent to

$$\delta_{UP} > \frac{(3n-1)}{(2n-1-\alpha)} \alpha. \tag{B8}$$

It is easy to verify that (B8) holds for $\alpha < 1/2$ if $n = 2$ and $\alpha < 1/4$ for $n = 3$.

[10] Expected welfare is higher if under pay-as-bid pricing if and only if $(1 - \tilde{\mu}^E)$ times the left-hand side is of (B7) is positive with $\mu = \tilde{\mu}^E$.

References

Ambrose, J. (2025). Two power station owners to get more than £12m for three hours of electricity. *The Guardian* .

Anatolitis, V. and Welisch, M. (2017). Putting renewable energy auctions into action – an agent-based model of onshore wind power auctions in germany. *Energy Policy* 110: 394–402, doi:10.1016/j.enpol.2017.08.024.

Anderson, E. and Holmberg, P. (2023). Multi-unit auctions with uncertain supply and single-unit demand. *EPRG Working Paper* doi:10.17863/CAM.105498.

Back, K. and Zender, J. F. (2001). Auctions of divisible goods with endogenous supply. *Economics Letters* 73: 29–34, doi:10.1016/s0165-1765(01)00465-7.

Bougt, D., Ghosh, G. and Liu, H. (2025). Revenue effects of ambiguity in multi-unit auctions. *Journal of Economic Theory* 225: 105996, doi:10.1016/j.jet.2025.105996.

Bower, J. and Bunn, D. (2001). Experimental analysis of the efficiency of uniform-price versus discriminatory auctions in the england and wales electricity market. *Journal of Economic Dynamics and Control* 25: 561–592, doi:10.1016/s0165-1889(00)00036-1.

Burkett, J. and Woodward, K. (2020). Uniform price auctions with a last accepted bid pricing rule. *Journal of Economic Theory* 185: 104954, doi:10.1016/j.jet.2019.104954.

Department for Energy Security & Net Zero (2023a). Electricity Generation Costs 2023. Tech. rep., Department for Energy Security & Net Zero.

Department for Energy Security & Net Zero (2023b). Review of Electricity Market Arrangements - Summary of responses to consultation. Tech. rep., Department for Energy Security & Net Zero.

EPEX (2025). Monthly power trading results of january 2025. https://www.epexspot.com/sites/default/files/download_center_files/2025-02-05_EPEX%20SPOT_Power%20Trading%20Results%2001-2025_0.pdf.

European Commission (2022). State of the union address 2022. https://state-of-the-union.ec.europa.eu/document/download/04c3ccc9-1237-4029-ae0e-

039709a0fb53_en?filename=SOTEU_2022_Address_EN.pdf.

Fabra, N., Fehr, N. von der and Harbord, D. (2006). Designing electricity auctions. *The RAND Journal of Economics* 37: 23–46, doi:10.1111/j.1756-2171.2006.tb00002.x.

Federico, G. and Rahman, D. (2003). Bidding in an electricity pay-as-bid auction. *Journal of Regulatory Economics* 24: 175–211, doi:10.1023/a:1024738128115.

Genc, T. S. (2008). Discriminatory versus uniform-price electricity auctions with supply function equilibrium. *Journal of Optimization Theory and Applications* 140: 9–31, doi:10.1007/s10957-008-9437-8.

Guerci, E., Ivaldi, S., Raberto, M. and Cincotti, S. (2007). Learning oligopolistic competition in electrcity auctions. *Computational Intelligence* 23: 197–220, doi:10.1111/j.1467-8640.2007.00298.x.

Holmberg, P. (2008). Unique supply function equilibrium with capacity constraints. *Energy Economics* 30: 148–172, doi:10.1016/j.eneco.2006.10.005.

Holmberg, P. (2009). Supply function equilibria of pay-as-bid auctions. *Journal of Regulatory Economics* 36: 154–177, doi:10.1007/s11149-009-9091-6.

Hästö, P. and Holmberg, P. (2006). Some inequalities related to the analysis of electricity auctions. *Applied Mathematics Letters* 19: 814–819, doi:10.1016/j.aml.2005.09.006.

International Energy Agency (2020). Projected Costs of Generating Electricity - 2020 Edition. Tech. rep., International Energy Agency.

Kamat, R. and Oren, S. S. (2002). Rational buyer meets rational seller: Reserves market equilibria under alternative auction designs. *Journal of Regulatory Economics* 21: 247–288, doi:10.1023/a:1015378914995.

Klemperer, P. D. and Meyer, M. A. (1989). Supply function equilibria in oligopoly under uncertainty. *Econometrica* 57: 1243–1277.

Krishna, V. (2010). *Auction theory*. Burlington, MA: Academic Press/Elsevier, 2nd ed., includes bibliographical references and index.

Pycia, M. and Woodward, K. (2025). A case for pay-as-bid auctions. *Journal of*

Political Economcy, forthcoming available on `https://sites.econ.uzh.ch/pycia/pycia-woodward-paba.pdf`, last accessed on 03/14/2024.

Rassenti, S. J., Smith, V. L. and Wilson, B. J. (2003). Discriminatory price auctions in electricity markets - low volatility at the expense of high price levels. *Journal of Regulatory Economics* 23: 109–123, doi:10.1023/a:1022250812631.

Sugianto, L. F. and Liao, K. Z. (2014). Comparison of different auction pricing rules in the electricity market. *Modern Applied Science* 8, doi:10.5539/mas.v8n1p147.

Viehmann, J., Lorenczik, S. and Malischek, R. (2021). Multi-unit multiple bid auctions in balancing markets: An agent-based q-learning approach. *Energy Economics* 93: 105035, doi:10.1016/j.eneco.2020.105035.

Xiong, G., Okuma, S. and Fujita, H. (2004). Multi-agent based experiments on uniform price and pay-as-bid electricity auction markets. In *2004 IEEE International Conference on Electric Utility Deregulation, Restructuring and Power Technologies. Proceedings*, DRPT-04. IEEE, 72–76, doi:10.1109/drpt.2004.1338471.

References

Abeler, J., Falk, A., Goette, L. and Huffman, D. (2011). Reference points and effort provision. *American Economic Review*, **101** (2), 470–492, doi: 10.1257/aer.101.2.470.

Akerlof, G. A. (1970). The market for lemons: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, **84** (3), 488–500, doi: 10.2307/1879431.

Allen, J., Clark, R., Hickman, B. and Richert, E. (2024). Resolving failed banks: Uncertainty, multiple bidding and auction design. *Review of Economic Studies*, **91** (3), 1201–1242, doi: 10.1093/restud/rdad062.

Ambrose, J. (2025). Two power station owners to get more than £12m for three hours of electricity. *The Guardian*, available on `https://www.theguardian.com/business/2025/jan/08/two-power-station-owners-to-get-more-than-12m-for-three-hours-of-\electricity`, last accessed on 15/12/2025.

Anatolitis, V. and Welisch, M. (2017). Putting renewable energy auctions into action an agent-based model of onshore wind power auctions in germany. *Energy Policy*, **110**, 394–402, doi: 10.1016/j.enpol.2017.08.024.

Anderson, E. and Holmberg, P. (2023). Multi-unit auctions with uncertain supply and single-unit demand. *EPRG Working Paper*, doi: 10.17863/CAM.105498, no. 2310.

Aussel, D., Bendotti, P. and Pitk, M. (2017a). Nash equilibrium in a pay-as-bid electricity market: Part 1 existence and characterization. *Optimization*, **66** (6), 1013–1025, doi: 10.1080/02331934.2016.1227981.

—, — and — (2017b). Nash equilibrium in a pay-as-bid electricity market part 2 - best response of a producer. *Optimization*, **66** (6), 1027–1053, doi: 10.1080/02331934.2016.1227982.

Auster, S. (2018). Robust contracting under common value uncertainty: Robust contracting. *Theoretical Economics*, **13** (1), 175–204, doi: 10.3982/te2385.

Back, K. and Zender, J. F. (2001). Auctions of divisible goods with endogenous supply. *Economics Letters*, **73** (1), 29–34, doi: 10.1016/s0165-1765(01)00465-7.

Bakirtzis, A. and Tellidou, A. (2006). Agent-based simulation of power markets under uniform and pay-as-bid pricing rules using reinforcement learning. In *2006 IEEE PES Power Systems Conference and Exposition*, IEEE, pp. 1168–1173.

BALZER, B. and ROSATO, A. (2025). Never say never: Optimal exclusion and reserve prices with expectations-based loss-averse buyers. *Journal of Economic Theory*, **228**, 106045, doi: 10.1016/j.jet.2025.106045.

—, — and VON WANGENHEIM, J. (2022). Dutch vs. first-price auctions with expectations-based loss-averse bidders. *Journal of Economic Theory*, **205**, 105545, doi: 10.1016/j.jet.2022.105545.

BARBERIS, N., HUANG, M. and SANTOS, T. (2001). Prospect theory and asset prices. *The Quarterly Journal of Economics*, **116** (1), 1–53, doi: 10.1162/003355301556310.

BAUCELLS, M., GRUSHKA-COCKAYNE, Y. and HWANG, W. (2024). Managerial mental accounting and downstream project decisions. *Management Science*, **70** (12), 8612–8630, doi: 10.1287/mnsc.2021.02929.

BERGEMANN, D. and HÖRNER, J. (2018). Should first-price auctions be transparent? *American Economic Journal: Microeconomics*, **10** (3), 177–218, doi: 10.1257/mic.20160278.

— and SCHLAG, K. H. (2008). Pricing without priors. *Journal of the European Economic Association*, **6** (23), 560–569, doi: 10.1162/jeea.2008.6.2-3.560.

— and — (2011). Robust monopoly pricing. *Journal of Economic Theory*, **146** (6), 2527–2543, doi: 10.1016/j.jet.2011.10.018.

BETTIS, R. A. (2017). Organizationally intractable decision problems and the intellectual virtues of heuristics. *Journal of Management*, **43** (8), 2620–2637, doi: 10.1177/0149206316679253.

BORENSTEIN, S. and BUSHNELL, J. (1999). An empirical analysis of the potential for market power in californias electricity industry. *The Journal of Industrial Economics*, **47** (3), 285–323, doi: 10.1111/1467-6451.00102.

—, BUSHNELL, J. B. and WOLAK, F. A. (2002). Measuring market inefficiencies in californias restructured wholesale electricity market. *American Economic Review*, **92** (5), 1376–1405, doi: 10.1257/000282802762024557.

BOSE, S., OZDENOREN, E. and PAPE, A. (2006). Optimal auctions with ambiguity. *Theoretical Economics*, **1** (4), 411–438.

BOUGT, D., GHOSH, G. and LIU, H. (2025). Revenue effects of ambiguity in multi-unit auctions. *Journal of Economic Theory*, **225**, 105996, doi: 10.1016/j.jet.2025.105996.

BOWER, J. and BUNN, D. (2001). Experimental analysis of the efficiency of uniform-price versus discriminatory auctions in the england and wales electricity market. *Journal of Economic Dynamics and Control*, **25** (34), 561–592, doi: 10.1016/s0165-1889(00)00036-1.

BURKETT, J. and WOODWARD, K. (2020a). Reserve prices eliminate low revenue equilibria in uniform price auctions. *Games and Economic Behavior*, **121**, 297–306, doi: 10.1016/j.geb.2020.02.011.

— and — (2020b). Uniform price auctions with a last accepted bid pricing rule. *Journal of Economic Theory*, **185**, 104954, doi: 10.1016/j.jet.2019.104954.

Bushnell, J. B., Mansur, E. T. and Saravia, C. (2008). Vertical arrangements, market structure, and competition: An analysis of restructured us electricity markets. *American Economic Review*, **98** (1), 237–266, doi: 10.1257/aer.98.1.237.

Börgers, T. (2015). *An Introduction to the Theory of Mechanism Design*. Oxford, New York: Oxford University Press.

—, Li, J. and Wang, K. (2025). Undominated mechanisms. *SSRN Working Paper*, doi: 10.2139/ssrn.5039409.

Camboni, R., Corazzini, L., Galavotti, S. and Valbonesi, P. (2025). Bidding on price and quality: An experiment on the complexity of scoring rule auctions. *Review of Economics and Statistics*, **107** (3), 755–770, doi: 10.1162/rest_a_01288.

Camerer, C. F. and Malmendier, U. (2007). Behavioral economics of organizations. In P. Diamond and H. Vartiainen (eds.), *Behavioral Economics and Its Applications*, *7*, Princeton University Press, pp. 235–290.

Caputo, M. R. (1990). Comparative dynamics via envelope methods in variational calculus. *The Review of Economic Studies*, **57** (4), 689–697, doi: 10.2307/2298093.

— (2005). *Foundations of Dynamic Economic Analysis: Optimal Control Theory and Applications*. Cambridge University Press.

Caragiannis, I., Jiang, Z. and Skoulakis, S. (2025). Rethinking pricing in energy markets: Pay-as-bid vs pay-as-clear. *Working Paper*, doi: 10.48550/ARXIV.2507.06035.

Carrasco, V., Farinha Luz, V., Kos, N., Messner, M., Monteiro, P. and Moreira, H. (2018). Optimal selling mechanisms under moment conditions. *Journal of Economic Theory*, **177**, 245–279, doi: 10.1016/j.jet.2018.05.005.

—, —, Monteiro, P. K. and Moreira, H. (2019). Robust mechanisms: the curvature case. *Economic Theory*, **68** (1), 203–222, doi: 10.1007/s00199-018-1120-1.

Carroll, G. (2017). Robustness and separation in multidimensional screening. *Econometrica*, **85** (2), 453–488, doi: 10.3982/ecta14165.

— (2019). Robustness in mechanism design and contracting. *Annual Review of Economics*, **11** (1), 139–166, doi: 10.1146/annurev-economics-080218-025616.

Cassady, R. (1967/1980). *Auctions and auctioneering*. Berkley and Los Angeles: University of California Press, reprint edn.

Che, Y.-K. and Zhong, W. (2024). Robustly optimal mechanisms for selling multiple goods. *Review of Economic Studies*, **92** (5), 2923–2951, doi: 10.1093/restud/rdae107.

Chiang, A. C. (1992). *Elements of Dynamic Optimization*. Waveland Press.

Clarke, E. H. (1971). Multipart pricing of public goods. *Public Choice*, **11** (1), 17–33, doi: 10.1007/bf01726210.

Cramton, P. (2004). Competitive bidding behavior in uniform-price auction markets. In *37th Annual Hawaii International Conference on System Sciences, 2004. Proceedings of the*, IEEE, pp. 1–11.

— and Ockenfels, A. (2017). The German 4G spectrum auction: Design and behaviour. *The Economic Journal*, **127** (605), F305–F324, doi: 10.1111/ecoj.12406.

— and Stoft, S. (2007). Why we need to stick with uniform-price auctions in electricity markets. *The Electricity Journal*, **20** (1), 26–37, doi: 10.1016/j.tej.2006.11.011.

Damianov, D. S. (2005). The uniform price auction with endogenous supply. *Economics Letters*, **88** (2), 152–158, doi: 10.1016/j.econlet.2005.02.004.

Dasgupta, S. and Spulber, D. F. (1989). Managing procurement auctions. *Information Economics and Policy*, **4** (1), 5–29, doi: 10.1016/0167-6245(89)90030-9.

Dastidar, K. G. (2008). On procurement auctions with fixed budgets. *Research in Economics*, **62** (2), 72–91, doi: 10.1016/j.rie.2008.01.001.

Dechenaux, E. and Kovenock, D. (2007). Tacit collusion and capacity withholding in repeated uniform price auctions. *The RAND Journal of Economics*, **38** (4), 1044–1069, doi: 10.1111/j.0741-6261.2007.00125.x.

Deck, C. A. and Wilson, B. J. (2008). Fixed revenue auctions: Theory and behavior. *Economic Inquiry*, **46** (3), 342–354, doi: 10.1111/j.1465-7295.2007.00090.x.

Department for Energy Security & Net Zero (2023a). Electricity generation costs 2023. Available on `https://assets.publishing.service.gov.uk/media/6556027d046ed400148b99fe/electricity-generation-costs-2023.pdf`, last accessed on 10/11/2025.

Department for Energy Security & Net Zero (2023b). Review of electricity market arrangements - summary of responses to consultation. Available on `https://assets.publishing.service.gov.uk/media/640226048fa8f527fe30dbba/review_of_electricity_market_arrangements_summary_of_responses.pdf`, last accessed on 10/11/2025.

Doraszelski, U., Lewis, G. and Pakes, A. (2018). Just starting out: Learning and equilibrium in a new market. *American Economic Review*, **108** (3), 565–615, doi: 10.1257/aer.20160177.

EPEX (2025). Monthly power trading results of january 2025. Available on `https://www.epexspot.com/sites/default/files/download_center_files/2025-02-05_EPEX%20SPOT_Power%20Trading%20Results%2001-2025_0.pdf`, last accessed on 21/03/2025.

Esö, P. and Futó, G. (1999). Auction design with a risk averse seller. *Economics Letters*, **65** (1), 71–74, doi: 10.1016/s0165-1765(99)00115-9.

European Commission (2022). State of the union address 2022. Available on `https://state-of-the-union.ec.europa.eu/document/download/04c3ccc9-1237-4029-ae0e-039709a0fb53_en?filename=SOTEU_2022_Address_EN.pdf`, last accessed on 21/03/2023.

European Commission (2024). Competitive bidding – a new tool for funding innovative low-carbon technologies under the innovation fund. Available on `https://climate.ec.europa.eu/eu-action/eu-funding-climate-action/innovation-fund/competitive-bidding_en`, last accessed on 07/19/2024.

Fabra, N. (2003). Tacit collusion in repeated auctions: Uniform versus discriminatory. *The Journal of Industrial Economics*, **51** (3), 271–293, doi: 10.1111/1467-6451.00201.

— and Llobet, G. (2023). Auctions with privately known capacities: Understanding competition among renewables. *The Economic Journal*, **133** (651), 1106–1146, doi: 10.1093/ej/ueac080.

—, von der Fehr, N. and Harbord, D. (2006). Designing electricity auctions. *The RAND Journal of Economics*, **37** (1), 23–46, doi: 10.1111/j.1756-2171.2006.tb00002.x.

—, von der Fehr, N. M. and de Frutos, M. (2011). Market design and investment incentives. *The Economic Journal*, **121** (557), 1340–1360, doi: 10.1111/j.1468-0297.2011.02450.x.

Federico, G. and Rahman, D. (2003). Bidding in an electricity pay-as-bid auction. *Journal of Regulatory Economics*, **24** (2), 175–211, doi: 10.1023/a:1024738128115.

Filiz-Ozbay, E. and Ozbay, E. Y. (2007). Auctions with anticipated regret: Theory and experiment. *American Economic Review*, **97** (4), 1407–1418, doi: 10.1257/aer.97.4.1407.

Galgana, R. and Golrezaei, N. (2025). Learning in repeated multiunit pay-as-bid auctions. *Manufacturing & Service Operations Management*, **27** (1), 200–229, doi: 10.1287/msom.2023.0403.

Gao, B. and Guo, L. (2024). Optimal contracts for time-inconsistent consumers with heterogeneous beliefs. *Management Science*, **71** (1), 861–878, doi: 10.1287/mnsc.2022.02653.

Genc, T. S. (2008). Discriminatory versus uniform-price electricity auctions with supply function equilibrium. *Journal of Optimization Theory and Applications*, **140** (1), 9–31, doi: 10.1007/s10957-008-9437-8.

Gibbard, A. (1973). Manipulation of voting schemes: A general result. *Econometrica*, **41** (4), 587–601, doi: 10.2307/1914083.

Graham, J. R., Harvey, C. R. and Puri, M. (2015). Capital allocation and delegation of decision-making authority within firms. *Journal of Financial Economics*, **115** (3), 449–470, doi: 10.1016/j.jfineco.2014.10.011.

Groves, T. (1973). Incentives in teams. *Econometrica*, **41** (4), 617–631, doi: 10.2307/1914085.

GUERCI, E., IVALDI, S., RABERTO, M. and CINCOTTI, S. (2007). Learning oligopolistic competition in electrcity auctions. *Computational Intelligence*, **23** (2), 197–220, doi: 10.1111/j.1467-8640.2007.00298.x.

— and RASTEGAR, M. A. (2012). Comparing system-marginal-price versus pay-as-bid auctions in a realistic electricity market scenario. In A. Teglio, S. Alfarano, E. Camacho-Cuena and M. Ginés-Vilar (eds.), *Managing Market Complexity*, *12*, Springer Berlin Heidelberg, pp. 141–153.

HAILU, A. and THOYER, S. (2007). Designing multiunit multiple bid auctions: An agentbased computational model of uniform, discriminatory and generalised vickrey auctions. *Economic Record*, **83**, S57–S72, doi: 10.1111/j.1475-4932.2007.00410.x.

HANSEN, R. G. (1988). Auctions with endogenous quantity. *The RAND Journal of Economics*, **19** (1), 44–58, doi: 10.2307/2555396.

HEIDHUES, P. and KŐSZEGI, B. (2008). Competition and price variation when consumers are loss averse. *American Economic Review*, **98** (4), 1245–1268, doi: 10.1257/aer.98.4.1245.

HEIM, S. and GÖTZ, G. (2021). Do pay-as-bid auctions favor collusion? evidence from germanys market for reserve power. *Energy Policy*, **155**, 112308, doi: 10.1016/j.enpol.2021.112308.

HIRTH, L., KHANNA, T. M. and RUHNAU, O. (2024). How aggregate electricity demand responds to hourly wholesale price fluctuations. *Energy Economics*, **135**, 107652, doi: 10.1016/j.eneco.2024.107652.

HOLMBERG, P. (2008). Unique supply function equilibrium with capacity constraints. *Energy Economics*, **30** (1), 148–172, doi: 10.1016/j.eneco.2006.10.005.

— (2009). Supply function equilibria of pay-as-bid auctions. *Journal of Regulatory Economics*, **36** (2), 154–177, doi: 10.1007/s11149-009-9091-6.

— and WOLAK, F. A. (2018). Comparing auction designs where suppliers have uncertain costs and uncertain pivotal status. *The RAND Journal of Economics*, **49** (4), 995–1027, doi: 10.1111/1756-2171.12259.

HORTAÇSU, A. and PULLER, S. L. (2008). Understanding strategic bidding in multiunit auctions: a case study of the texas electricity spot market. *The RAND Journal of Economics*, **39** (1), 86–114, doi: 10.1111/j.0741-6261.2008.00005.x.

HURWICZ, L. (1972). On informationally decentralized systems. In R. Radner and C. B. McGuire (eds.), *Decision and Organization: A Volume in Honor of Jacob Marschak*, Amsterdam: North-Holland Publishing, pp. 297–336.

HÄSTÖ, P. and HOLMBERG, P. (2006). Some inequalities related to the analysis of electricity auctions. *Applied Mathematics Letters*, **19** (8), 814–819, doi: 10.1016/j.aml.2005.09.006.

INTERNATIONAL ENERGY AGENCY (2020). Projected costs of generating electricity - 2020 edition. Available on `https://iea.blob.core.windows.net/assets/ae17da3d-e8a5-4163-a3ec-2e6fb0b5677d/Projected-Costs-of-Generating-Electricity-2020.pdf`, last accessed on 10/11/2025.

KAHNEMAN, D. and TVERSKY, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, **47** (2), 263–292, doi: 10.2307/1914185.

KAMAT, R. and OREN, S. S. (2002). Rational buyer meets rational seller: Reserves market equilibria under alternative auction designs. *Journal of Regulatory Economics*, **21** (3), 247–288, doi: 10.1023/a:1015378914995.

KASBERGER, B. and WOODWARD, K. (2025). Bidding in multi-unit auctions under limited information. *Journal of Economic Theory*, **226**, 106008, doi: 10.1016/j.jet.2025.106008.

KIM, H. (2025). The value of competitor information: Evidence from a field experiment. *Management Science*, **71** (4), 3600–3621, doi: 10.1287/mnsc.2022.04062.

KLEMPERER, P. D. and MEYER, M. A. (1989). Supply function equilibria in oligopoly under uncertainty. *Econometrica*, **57** (6), 1243–1277, doi: 10.2307/1913707.

KNIGHT, F. H. (1921). *Risk, Uncertainty and Profit*. Boston and New York: Houghton Mifflin Company.

KNITTEL, C. R. and STANGO, V. (2003). Price ceilings as focal points for tacit collusion: Evidence from credit cards. *American Economic Review*, **93** (5), 1703–1729, doi: 10.1257/000282803322655509.

KOÇYIIT, Ç., KUHN, D. and RUJEERAPAIBOON, N. (2024). Regret minimization and separation in multi-bidder, multi-item auctions. *INFORMS Journal on Computing*, **36** (6), 1543–1561, doi: 10.1287/ijoc.2022.0275.

KŐSZEGI, B. and RABIN, M. (2006). A model of reference-dependent preferences. *The Quarterly Journal of Economics*, **121** (4), 1133–1165, doi: 10.1093/qje/121.4.1133.

— and RABIN, M. (2007). Reference-dependent risk attitudes. *American Economic Review*, **97** (4), 1047–1073, doi: 10.1257/aer.97.4.1047.

— and RABIN, M. (2009). Reference-dependent consumption plans. *American Economic Review*, **99** (3), 909–936, doi: 10.1257/aer.99.3.909.

KREMER, I. and NYBORG, K. G. (2004). Underpricing and market power in uniform price auctions. *The Review of Financial Studies*, **17** (3), 849–877, doi: 10.1093/rfs/hhg051.

KRISHNA, V. (2010). *Auction theory*. Burlington, MA: Academic Press/Elsevier, 2nd edn.

KRZYWNICKA, M. and BARNER, L. (2025). On strategic behavior in electricity markets during the 2022 energy crisis: An ex-post analysis for the german power sector. *Energy Reports*, **13**, 5474–5485, doi: 10.1016/j.egyr.2025.04.040.

LABANDEIRA, X., LABEAGA, J. M. and LÓPEZ-OTERO, X. (2017). A meta-analysis on the price elasticity of energy demand. *Energy Policy*, **102**, 549–568, doi: 10.1016/j.enpol.2017.01.002.

LAFFONT, J.-J. and MARTIMORT, D. (2002). *The Theory of Incentives: The Principal-Agent Model*. Princeton, NJ: Princeton Univ. Press.

LAITSOS, V., VONTZOS, G., PARASCHOUDIS, P., TSAMPASIS, E., BARGIOTAS, D. and TSOUKALAS, L. H. (2024). The state of the art electricity load and price forecasting for the modern wholesale electricity market. *Energies*, **17** (22), 5797, doi: 10.3390/en17225797.

LANGE, A. and RATAN, A. (2010). Multi-dimensional reference-dependent preferences in sealed-bid auctions how (most) laboratory experiments differ from the field. *Games and Economic Behavior*, **68** (2), 634–645, doi: 10.1016/j.geb.2009.08.008.

LENGWILER, Y. (1999). The multiple unit auction with variable supply. *Economic Theory*, **14** (2), 373–392, doi: 10.1007/s001990050299.

LI, S. (2024). Designing simple mechanisms. *Journal of Economic Perspectives*, **38** (4), 175–192, doi: 10.1257/jep.38.4.175.

— and LO, I. (2023). New solution concepts. In N. I. Federico Echenique and V. V. Vazriani (eds.), *Online and Matching-Based Marketing Design*, *27*, Cambridge University Press, pp. 575 – 590.

LINNENBRINK, D. (2025). Pay-as-bid auctions with private information. *SSRN Working Paper*, doi: 10.2139/ssrn.4416398.

LIU, T. (2007). *Essays on Corporate Finance Using an Auction Approach*. Ph.D. thesis, Carnegie Mellon University.

— and PARLOUR, C. A. (2023). Raising money. *SSRN Working Paper*, doi: 10.2139/ssrn.4431934, available on `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4431934`, last accessed on 09/09/2025.

LIU, Z., YAN, J., SHI, Y., ZHU, K. and PU, G. (2012). Multi-agent based experimental analysis on bidding mechanism in electricity auction markets. *International Journal of Electrical Power & Energy Systems*, **43** (1), 696–702, doi: 10.1016/j.ijepes.2012.05.056.

LONG, X., NASIRY, J. and WU, Y. (2020). A behavioral study on abandonment decisions in multistage projects. *Management Science*, **66** (5), 1999–2016, doi: 10.1287/mnsc.2018.3270.

MACKAY, E. (2017). 10 of the most bizarre things sold at music auctions. *BBC*, available at `https://www.bbc.co.uk/music/articles/4e641c9b-1a1f-4b07-9201-5843854aaa25`, last accessed 09/09/2025.

MARZILLI ERICSON, K. M. and FUSTER, A. (2011). Expectations as endowments: Evidence on reference-dependent preferences from exchange and valuation experiments. *The Quarterly Journal of Economics*, **126** (4), 1879–1907, doi: 10.1093/qje/qjr034.

MAS-COLLEL, A., WHINSTON, M. D. and GREEN, J. R. (1995). *Microeconomic Theory*. Oxford, New York: Oxford University Press.

MCADAMS, D. (2007). Adjustable supply in uniform price auctions: Non-commitment as a strategic tool. *Economics Letters*, **95** (1), 48–53, doi: 10.1016/j.econlet.2006.08.035.

McNeal, R. A. (1988). The brides of babylon: Herodotus 1.196. *Historia: Zeitschrift für Alte Geschichte*, **37** (1), 54–71.

Myerson, R. B. (1979). Incentive compatibility and the bargaining problem. *Econometrica*, **47** (1), 61–73, doi: 10.2307/1912346.

— (1981). Optimal auction design. *Mathematics of Operations Research*, **6** (1), 58–73, doi: 10.1287/moor.6.1.58.

— (1982). Optimal coordination mechanisms in generalized principal-agent problems. *Journal of Mathematical Economics*, **10** (1), 67–81, doi: 10.1016/0304-4068(82)90006-4.

— (1986). Multistage games with communication. *Econometrica*, **54** (2), 323–358, doi: 10.2307/1913154.

Oniki, H. (1973). Comparative dynamics (sensitivity analysis) in optimal control theory. *Journal of Economic Theory*, **6** (3), 265–283, doi: 10.1016/0022-0531(73)90050-1.

Pycia, M. and Woodward, K. (2026). A case for pay-as-bid auctions. *Journal of Political Economy*, **134** (2), 795–845, doi: 10.1086/738341.

Rosato, A. and Tymula, A. A. (2019). Loss aversion and competition in vickrey auctions: Money aint no good. *Games and Economic Behavior*, **115**, 188–208, doi: 10.1016/j.geb.2019.02.014.

Rothkopf, M. H., Teisberg, T. J. and Kahn, E. P. (1990). Why are Vickrey auctions rare? *Journal of Political Economy*, **98** (1), 94–109, doi: 10.1086/261670.

Rothschild, M. and Stiglitz, J. (1976). Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, **90** (4), 629–649, doi: 10.2307/1885326.

RTR (2019). Wir bringen 5G in die Fläche. Available on `https://www.rtr.at/TKP/presse/pressemitteilungen/pressemitteilungen/PI23092019TK.de.html`, last accessed on 09/09/2025.

— (2020). 5G-Auktion macht nahezu flächendeckende Mobilfunk-Breitbandversorgung in Zukunft möglich. Available on `https://www.rtr.at/TKP/presse/pressemitteilungen/pressemitteilungen/pinfo11092020tkp.de.html`, last accessed on 21/11/2022.

Shalev, J. (2000). Loss aversion equilibrium. *International Journal of Game Theory*, **29** (2), 269–287, doi: 10.1007/s001820000038.

Shunda, N. (2009). Auctions with a buy price: The case of reference-dependent preferences. *Games and Economic Behavior*, **67** (2), 645–664, doi: 10.1016/j.geb.2009.04.013.

Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, **69** (1), 99–118, doi: 10.2307/1884852.

Son, Y., Baldick, R., Lee, K.-H. and Siddiqi, S. (2004). Short-term electricity market auction game analysis: Uniform and pay-as-bid pricing. *IEEE Transactions on Power Systems*, **19** (4), 1990–1998, doi: 10.1109/tpwrs.2004.836192.

Song, Y., Chu, Y., Yoon, Y. and Jin, Y. (2025). Virtual power plant bidding strategies in pay-as-bid and pay-as-clear markets: Analysis of imbalance penalties and market operations. *Energies*, **18** (6), 1383, doi: 10.3390/en18061383.

Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, **87** (3), 355–374, doi: 10.2307/1882010.

Sugianto, L. F. and Liao, K. Z. (2014). Comparison of different auction pricing rules in the electricity market. *Modern Applied Science*, **8** (1), doi: 10.5539/mas.v8n1p147.

Thaler, R. H. (1985). Mental accounting and consumer choice. *Marketing Science*, **4** (3), 199–214, doi: 10.1287/mksc.4.3.199.

— (1999). Mental accounting matters. *Journal of Behavioral Decision Making*, **12** (3), 183–206, doi: 10.1002/(sici)1099-0771(199909)12:3<183::aid-bdm318>3.0.co;2-f.

Tillio, A. d., Kos, N. and Messner, M. (2016). The design of ambiguous mechanisms. *The Review of Economic Studies*, **84** (1), 237–276, doi: 10.1093/restud/rdw051.

TKK (2019). Tender document in the procedure for awarding spectrum in the 700, 1500 and 2100 MHz ranges. Available on https://www.rtr.at/TKP/was_wir_tun/telekommunikation/spectrum/procedures/Multibandauktion_700-1500-2100MHz_2020/FRQ5G_2020_tender_document.de.html, last accessed on 16/12/2022.

Toikka, J. (2011). Ironing without control. *Journal of Economic Theory*, **146** (6), 2510–2526, doi: 10.1016/j.jet.2011.06.003.

Turcic, D., Markou, P., Kouvelis, P. and Corsten, D. (2024). Automotive procurement under opaque prices: Theory with evidence from the BMW supply chain. *Management Science*, **70** (6), 3664–3683, doi: 10.1287/mnsc.2023.4880.

UK Government (2024). Record breaking funding for clean energy in Britain. Available on https://www.gov.uk/government/news/record-breaking-funding-for-clean-energy-in-britain, last accessed 08/23/2024.

Vanelli, M., Como, G. and Fagnani, F. (2025). How competitive are pay-as-bid auction games? *Working Paper*, doi: 10.48550/ARXIV.2504.13920.

Vickrey, W. (1961). Counterspeculation, auctions and competitive sealed tenders. *The Journal of Finance*, **16** (1), 8–37, doi: 10.1111/j.1540-6261.1961.tb02789.x.

— (1962). Auction and bidding games. In *Recent Advances in Game Theory*, Princetion, New Jersey: Princeton University Conference, pp. 15–27.

Viehmann, J., Lorenczik, S. and Malischek, R. (2021). Multi-unit multiple bid auctions in balancing markets: An agent-based q-learning approach. *Energy Economics*, **93**, 105035, doi: 10.1016/j.eneco.2020.105035.

Weitzman, M. L. (1974). Prices vs. quantities. *The Review of Economic Studies*, **41** (4), 477–491, doi: 10.2307/2296698.

Willems, B. and Yueting, Y. (2023). Bidding and investment in wholesale electricy markets: Discriminatory versus uniform-price auctions. *Toulouse School of Economics Working Paper*, (1462).

Wilson, R. (1979). Auctions of shares. *The Quarterly Journal of Economics*, **93** (4), 675–689, doi: 10.2307/1884475.

Wolak, F. A. (2003). Identification and estimation of cost functions using observed bid data. In L. P. H. Mathias von Dewatripont and S. J. Turnovsky (eds.), *Advances in Economics and Econometrics – Theory and Applications, Eigth World Congress*, *4*, Cambridge University Press, pp. 133–169.

Wolitzky, A. (2016). Mechanism design with maxmin agents: Theory and an application to bilateral trade. *Theoretical Economics*, **11** (3), 971–1004, doi: 10.3982/te2089.

Xiong, G., Okuma, S. and Fujita, H. (2004). Multi-agent based experiments on uniform price and pay-as-bid electricity auction markets. In *2004 IEEE International Conference on Electric Utility Deregulation, Restructuring and Power Technologies. Proceedings*, DRPT-04, IEEE, pp. 72–76.

Zhang, W. (2022). Auctioning multiple goods without priors. *Working Paper*, doi: 10.48550/ARXIV.2204.13726.

Zhao, D., Botterud, A. and Ilic, M. (2023). Uniform pricing vs pay as bid in 100%-renewables electricity markets: A game-theoretical analysis. In *Proceedings of the 14th ACM International Conference on Future Energy Systems*, e-Energy 23, ACM, pp. 236–241.