# Essays on

# Digital Currencies and Monetary Policy

Dissertation

zur Erlangung des Grades eines Doktors der Wirtschaftswissenschaft

der Rechts- und Wirtschaftswissenschaftlichen

Fakultät der Universität Bayreuth

vorgelegt von

Jonas Groß

aus Bayreuth

| | |
|---|---|
| Dekan: | Prof. Dr. Jörg Schlüchtermann |
| Erstberichterstatter: | Prof. Dr. Bernhard Herz |
| Zweitberichterstatter: | Prof. Dr. Martin Leschke |
| Drittberichterstatter: | Prof. Dr. Nils Urbach |
| Termin der mündlichen Prüfung: | 31. Mai 2022 |

# Danksagung

An dieser Stelle möchte ich mich bei all den Personen bedanken, die mich bei meinem Promotionsvorhaben tatkräftig unterstützt haben.

Besonderer Dank gilt hierbei meinem Doktorvater Bernhard Herz, der meine Promotion überhaupt erst ermöglicht hat. Durch seine herzliche und hilfsbereite Betreuung habe ich während der Promotionszeit kontinuierlich dazugelernt. Sein konstruktives und ehrliches Feedback hat die Qualität meiner wissenschaftlichen Papiere maßgeblich verbessert. Außerordentlich dankbar bin ich dir, lieber Bernhard, dafür, dass du aufgrund deiner Offenheit hinsichtlich innovativer ökonomischer Themen und Thesen all meine Forschungsideen unterstützt hast. So war es mir eine Freude, mich im Rahmen meines Promotionsvorhabens mit Themen rund um Geldpolitik und digitale Währungen zu beschäftigen. Ihnen gilt meine besondere Leidenschaft. Zu der vertieften Auseinandersetzung mit diesen Themen hast du mir den Weg geebnet, als du mir Ende 2018 die Teilnahme an einer Konferenz zu digitalen Währungen an der Frankfurt School of Finance & Management nahegelegt hattest.

Außerdem möchte ich meinen herausragenden Ko-Autoren danken, durch deren Expertise ich bei jedem Forschungsprojekt meine fachlichen und methodischen Kenntnisse schärfen konnte. Mir werden vor allem die gemeinsamen Projekte mit meinen langjährigen Studienkommilitonen und guten Freunden Johannes Zahner und Jonathan Schiller in Erinnerung bleiben. Hier kam auch der Spaß nie zu kurz. Die geschätzten Kollegen am Lehrstuhl und im gesamten VWL-Fachbereich haben durch ihre Hilfsbereitschaft und Kompetenz ebenfalls zu meinem erfolgreichen Promotionsvorhaben beigetragen. Besonders möchte ich hier Sabine Übelhack hervorheben. Liebe Sabine, durch deine Herzlichkeit und Hilfsbereitschaft trägst du maßgeblich zu der offenen und positiven Arbeitsatmosphäre am Lehrstuhl bei.

Außerdem möchte ich mich noch bei meiner Familie bedanken. Ohne eure Unterstützung, euren Zuspruch, eure Hilfsbereitschaft und euer Feedback wäre ich nicht da, wo ich heute bin. Das kritische Hinterfragen ökonomischer Themen und Theorien habt ihr stets gefördert, wovon ich noch heute profitiere. Auch die vielen Diskussionen mit meinen Großeltern, vor allem rund um digitale Währungen, haben mir häufig eine alternative Sichtweise und Denkanstöße gegeben, die ich in meine Überlegungen und Papiere habe einfließen lassen können.

Besonderer Dank gilt zudem meiner Partnerin Stefanie Simon; einer unverzichtbaren Stütze. Steffi, du hast mich immer auf meinem Weg und bei all meinen Vorhaben unterstützt und mich auch in Phasen des Zweifelns aufgebaut. Ich weiß es sehr zu schätzen, dass du meine arbeitsintensiven Wochenenden und mein frühes Aufstehen toleriert hast.

# Abstract

This dissertation studies issues related to digital currencies and monetary policy. In particular, it analyzes the determinants of the monetary policy of the European Central Bank (ECB) and examines design aspects of central bank digital currencies (CBDCs), e.g., related to financial stability, monetary policy, and privacy.

Almost 15 years ago, the US bank Lehman Brothers went bankrupt. This incident marked the beginning of the global financial crisis — a crisis that shook the monetary system. Central banks worldwide introduced unprecedented monetary policy measures to stimulate inflation and economic activity. Since then, the environment for conducting monetary policy has changed due to the structural effects of the crisis. Assets prices, public debt, and the size of central banks' balance sheets increased substantially and reached all their all-time highs in the euro area in 2021.

Besides these economic effects, the global financial crisis was also the starting point of a monetary revolution. As a reaction to the crisis, Satoshi Nakamoto (Nakamoto, 2008) — a pseudonym whose identity remains unknown — initiated a novel decentralized payment system that operates without banks, central banks, and government interference: Bitcoin was born. Originating from Bitcoin, by December 2021, more than 10,000 such cryptocurrencies with a market capitalization of more than two trillion US dollars emerged (Coinmarketcap, 2021). As a con-

sequence of these monetary innovations and the ongoing digitization of payments, central banks consider issuing own digital currencies, so-called CBDCs.[1] CBDCs promise to combine the advantages of (physical) cash and (digital) bank deposits, e.g., cheap, convenient, and safe payments in the digital realm. However, they could also impair financial stability and undermine data privacy. As CBDCs have not been introduced in advanced economies yet, substantial theory-based analyses are required to study their benefits and risks adequately.

Chapter 1 addresses the question, which factors influenced the ECB's monetary policy in the period from 1999 to 2018. After identifying potential monetary policy determinants based on a literature review and textual analysis of the ECB's public communication, we use an empirical Bayesian model averaging (BMA) approach to determine key variables that impacted the ECB's monetary policy decisions before and after the global financial crisis. While in the literature, researchers typically select one model to study monetary policy determinants, using BMA allows accounting for uncertainty about the choice of the respective empirical model. This uncertainty arises, amongst others, from the heterogeneity in the ECB's decision-making body, the Governing Council. Our analysis considers approximately 33,000 different empirical models that can be constructed from the identified potential relevant determinants. Our results suggest the following: First, in the time period analyzed, the ECB mainly focused on the inflation rate when setting interest rates. Second, economic activity indicators were in the focus of the ECB before the financial crisis. Third, over the last decade, the role of economic activity decreased, thereby supporting the hypothesis that inflation was the main driver of monetary policy decisions in the post-crisis period. This result is supported by findings from textual analysis. These results show that, in recent years, official ECB communication mentioned terms related to inflation more frequently than terms related to

[1] This dissertation focuses on retail CBDCs, i.e., digital currencies issued by central banks available for private agents. Wholesale CBDCs, which are only available for selected entities, such as banks, are excluded from the analysis.

economic activity. Fourth, when setting interest rates, central bankers appeared to consider more than one model, favoring the use of averaging techniques for studying monetary policy determinants.

Chapter 2 focuses on CBDCs. It studies the effects of CBDCs on the financial sector and monetary policy, and analyzes which measures the central bank can undertake to prevent destabilizing effects for the economy. While CBDCs might offer several benefits for their users, they potentially impose threats to the financial sector. They could disintermediate commercial banks and facilitate bank runs since CBDCs, in contrast to commercial bank money, constitute digital forms of central bank money with marginal risk. Thus, in times of crises, private agents could decide to convert substantial amounts of commercial bank money in CBDC, thereby posing a risk to banks' liquidity. To analyze these concerns in the absence of any CBDC-specific empirical data, we develop a New Keynesian dynamic stochastic general equilibrium model and simulate a financial crisis in a world with and without CBDC. In particular, we compare the effects of remunerated (interest-bearing) and non-remunerated (non-interest-bearing) CBDCs. We find that CBDCs indeed crowd out bank deposits and negatively affect bank funding. However, this crowding-out effect can be mitigated if the central bank chooses to provide additional central bank funds or to disincentivize large-scale CBDC accumulation via low or potentially even negative interest rates. Thus, our results suggest that a CBDC does not necessarily impair the financial sector if the central bank chooses adequate design and policy measures.

Chapter 3 studies the design of CBDCs in more detail and addresses the question of how a CBDC should be designed to ensure a high degree of data privacy while taking legal requirements into account. While physical cash allows anonymous transactions, i.e., transaction data is only observable for the two transaction parties involved and not for third parties, such as banks and central banks, CBDC data is generally recorded in a digital database. Opponents of CBDCs criticize that data access for the central bank could incentivize the misuse of CBDC payment data by public sector entities and could undermine citizens' data privacy. In Chapter 3, utilizing a design science research approach, we propose a CBDC system that preserves citizens' payment privacy. Despite being often perceived as a tradeoff, we show that it is feasible to provide anonymity for CBDC payments while also complying with regulations related to anti-money laundering and combating the financing of terrorism. Privacy and compliance are guaranteed by zero-knowledge proofs, cryptographic innovations.

# Zusammenfassung

Die vorliegende Dissertation behandelt Themen rund um digitale Währungen und Geldpolitik. Es werden insbesondere die Determinanten der Geldpolitik der Europäischen Zentralbank (EZB) analysiert und Ausgestaltungsfragen von digitalen Zentralbankwährungen (CBDCs) untersucht, z.B. in Bezug auf Finanzstabilität, Geldpolitik und Datenschutz.

Vor fast 15 Jahren musste die US-Bank Lehman Brothers Insolvenz anmelden. Dieses Ereignis markierte den Beginn der globalen Finanzkrise — eine Krise, die das Geldsystem erschütterte. Zentralbanken starteten auf der ganzen Welt beispiellose expansive geldpolitische Maßnahmen, um die Inflation und die wirtschaftliche Aktivität anzukurbeln. Seitdem hat sich das geldpolitische Umfeld aufgrund struktureller Auswirkungen der Krise nachhaltig verändert. Vermögenspreise, Staatsverschuldungen und Bilanzsummen der Notenbanken sind erheblich gestiegen und haben im Jahre 2021 allesamt in der Eurozone Höchststände erreicht.

Neben den wirtschaftlichen Auswirkungen war die globale Finanzkrise auch der Beginn einer monetären Revolution. Als Reaktion auf die Krise initiierte Satoshi Nakamoto (Nakamoto, 2008) — ein Pseudonym mit bis heute unbekannter Identität — ein neuartiges, dezentralisiertes Geldsystem, das ohne Banken, Zentralbanken und staatliche Behörden auskommt: Der Bitcoin wurde geboren. Im Dezember 2021

gab es mehr als 10.000 solcher Kryptowährungen mit einer Marktkapitalisierung von mehr als zwei Billionen US-Dollar (Coinmarketcap, 2021).

Aufgrund dieser monetären Innovationen sowie der fortschreitenden Digitalisierung des Zahlungsverkehrs prüfen Zentralbanken derzeit intensiv die Ausgabe eigener digitaler Währungen.[2] CBDCs versprechen die Vorteile des (physischen) Bargelds und des (digitalen) Giralgelds zu vereinen. So sollen CBDCs günstige, nutzerfreundliche, bequeme und sichere digitale Zahlungen ermöglichen. Allerdings könnten CBDCs im Extremfall auch die Finanzstabilität gefährden und den Datenschutz untergraben. Da CBDCs in fortgeschrittenen Volkswirtschaften bislang noch nicht eingeführt wurden, sind theoriebasierte Analysen und Simulationen notwendig, um die Vorteile und Risiken von CBDCs verstehen und evaluieren zu können.

Kapitel 1 widmet sich der Frage, welche ökonomischen Kenngrößen die Geldpolitik der EZB in der Zeit von 1999 bis 2018 beeinflusst haben. Im Rahmen der Analyse werden zunächst potenzielle Einflussfaktoren mit Hilfe einer Literaturrecherche und Textanalyse der Kommunikation der EZB identifiziert. Im Anschluss werden — basierend auf einem empirischen Bayesian Model Averaging (BMA)-Ansatz — diejenigen Einflussfaktoren abgeleitet, welche die geldpolitischen Entscheidungen der EZB vor und nach der globalen Finanzkrise beeinflusst haben. Während in der bestehenden Literatur üblicherweise ein einziges Modell zur Untersuchung herangezogen wird, ermöglicht es ein BMA-Ansatz, eine Vielzahl von Modellen zu berücksichtigen. Folglich können Unsicherheiten hinsichtlich der Wahl des empirischen Modells berücksichtigt werden, die sich u.a. aus der heterogenen Zusammensetzung des Entscheidungsgremiums der EZB, des EZB-Rats, ergeben. Die Analyse berücksichtigt ca. 33.000 verschiedene empirische Modelle, die aus den identifizierten, potenziell relevanten, Einflussfaktoren konstruiert werden. Unsere Kernergebnisse lassen sich wie folgt zusammenfassen: Erstens orientierte sich die

[2]Die vorliegende Dissertation beschäftigt sich mit Retail CBDCs, d. h. digitalen Zentralbankwährungen, die für die Allgemeinheit verfügbar sind. Wholesale CBDCs, die nur für bestimmte Entitäten, wie z. B. Banken, verfügbar sind, werden in der Analyse nicht berücksichtigt.

EZB im Untersuchungszeitraum bei ihren Zinsentscheidungen hauptsächlich an Daten zur Inflationsrate. Zweitens standen vor der Finanzkrise auch Indikatoren der wirtschaftlichen Aktivität im Fokus der EZB. Drittens nahm die Bedeutung der wirtschaftlichen Aktivität im letzten Jahrzehnt ab, was ebenfalls die Hypothese stützt, dass die Inflation in der Zeit nach der Krise der wichtigste Einflussfaktor auf geldpolitische Entscheidungen war. Diese Schlussfolgerung wird durch Ergebnisse der Textanalyse untermauert, welche zeigen, dass in der offiziellen Kommunikation der EZB im letzten Jahrzehnt Begriffe im Zusammenhang mit Inflation häufiger genannt wurden als Begriffe im Zusammenhang mit wirtschaftlicher Aktivität. Viertens schienen die Notenbanker bei der Festlegung der Zinssätze mehr als ein Modell in Betracht zu ziehen, weshalb Averaging-Verfahren in der Forschung zur Untersuchung geldpolitischer Determinanten eine noch prominentere Rolle einnehmen sollten.

Kapitel 2 beschäftigt sich mit CBDCs und untersucht, welche Auswirkungen eine CBDC auf den Finanzsektor und die Geldpolitik der Zentralbanken haben kann. Zudem wird aufgezeigt, welche Maßnahmen die Zentralbank ergreifen kann, um destabilisierende Effekte für die Ökonomie zu verhindern. Während CBDCs für Nutzer zahlreiche Vorteile bieten können, stellen sie für den Finanzsektor eine potenzielle Gefahr dar. CBDCs könnten Geschäftsbanken aus dem Markt drängen, eine sogenannte Disintermediation des Finanzsektors vorantreiben und Bank Runs begünstigen, da es sich bei CBDCs — im Gegensatz zu Giralgeld — um eine digitale Form von Zentralbankgeld mit geringem Risiko handelt. In Krisenzeiten könnten Kunden daher beschließen, erhebliche Beträge Giralgeld in CBDCs umzutauschen, was ein Risiko für die Liquidität der Banken darstellen könnte. Um diese Bedenken adäquat — und ohne die Verfügbarkeit empirischer Daten zu CBDCs — zu analysieren, wurde ein neukeynesianisches dynamisches stochastisches allgemeines Gleichgewichtsmodell entwickelt und eine Finanzkrise in einer Welt mit und ohne CBDC simuliert. Insbesondere wurden hierbei die Effekte von verzinsten und nicht verzinsten CBDCs verglichen. Die Analyse kommt zu dem Ergebnis, dass CBDCs tatsächlich Bankeinlagen verdrängen und die Refinanzierung von Banken negativ beeinflussen können. Dieser Verdrängungseffekt kann jedoch abgeschwächt werden, wenn die Zentralbank beschließen würde, zusätzliches Zentralbankgeld für Banken zur Verfügung zu stellen oder die Akkumulation von großen CBDC-Beständen durch niedrige oder sogar negative Zinssätze zu verhindern. Die Ergebnisse deuten demnach darauf hin, dass eine CBDC den Finanzsektor nicht notwendigerweise negativ beeinträchtigen würde, wenn die Zentralbank geeignete Gegenmaßnahmen ergreifen würde.

Kapitel 3 widmet sich der Frage, wie eine CBDC ausgestaltet sein sollte, um einen hohen Grad an Datenschutz gewährleisten und rechtliche Vorgaben berücksichtigen zu können. Während Bargeld heute anonyme Transaktionen ermöglicht, bei denen Transaktionsdaten nur für die beteiligten Parteien und nicht für Dritte, wie Banken und Zentralbanken, einsehbar sind, werden CBDC-Daten im Allgemeinen in einer digitalen Datenbank aufgezeichnet. Kritiker befürchten, dass der Zugriff der Zentralbank auf vertrauliche Zahlungsdaten Anreize zum Datenmissbrauch durch öffentliche Stellen schaffen und einen adäquaten Datenschutz der Bürger untergraben könnte. In Kapitel 3 wird anhand eines Design Science Research-Ansatzes skizziert, wie ein CBDC-System ausgestaltet werden kann, das die Privatsphäre der Bürger schützt. Während bislang häufig von einem Trade-off ausgegangen wird, zeigt sich, dass es möglich ist, Anonymität für Zahlungen zu gewährleisten und gleichzeitig rechtliche Vorschriften rund um Geldwäsche und Terrorfinanzierung einzuhalten. Dieser Schutz der Privatsphäre und die Einhaltung rechtlicher Vorschriften werden durch Zero-Knowledge-Proofs — kryptografische Innovationen — gewährleistet.

# Contents

# List of Figures

# List of Tables

# List of Acronyms

| | |
|---|---|
| **AIC** | Akaike information criterion |
| **AML** | Anti-money laundering |
| **ATM** | Automated teller machine |
| **BIC** | Bayesian information criterion |
| **BMA** | Bayesian model averaging |
| **BRIC** | Benchmark risk inflation criterion |
| **CBDC** | Central bank digital currency |
| **CES** | Constant elasticity of substitution |
| **CFT** | Combating the financing of terrorism |
| **CISS** | Composite Index of Systemic Stress |
| **DSGE** | Dynamic stochastic general equilibrium |
| **DSR** | Design science research |
| **ECB** | European Central Bank |
| **EONIA** | Euro OverNight Index Average |
| **ELB** | Effective lower bound |
| **EU** | European Union |
| **Fed** | Federal Reserve |
| **GDP** | Gross domestic product |
| **HICP** | Harmonized Index of Consumer Prices |
| **HP** | Hodrick-Prescott |
| **IRFs** | Impulse response functions |
| **IS** | Information systems |
| **KYC** | Know-your-customer |
| **MSE** | Mean squared error |
| **MRO** | Main refinancing operations |
| **OLS** | Ordinary least squares |
| **PIP** | Posterior inclusion probability |
| **PSP** | Payment service provider |
| **RIC** | Risk inflation criterion |
| **RTD** | Real-Time Database |
| **SNARK** | Zero-knowledge succinct non-interactive argument of knowledge |
| **SPF** | Survey of Professional Forecasters |
| **UAS** | Unspent account state |
| **UIP** | Unit information prior |
| **UTXO** | Unspent transaction output |
| **ZKP** | Zero-knowledge proof |

# Chapter 1

# What Is on the ECB's Mind? Monetary Policy Before and After the Global Financial Crisis

**Abstract**

This paper analyzes the monetary policy of the European Central Bank (ECB) both before and after the outbreak of the global financial crisis in 2008. In the literature, researchers typically select *one* Taylor rule-based model to analyze monetary policy of central banks and to derive determinants for the interest rate setting. However, uncertainty about the choice of this respective model is typically neglected. In contrast, we apply a Bayesian model averaging (BMA) approach to extend the Taylor rule to account for model uncertainty driven by heterogeneity in the ECB's decision-making body, the Governing Council. Our results suggest the following: First, the ECB focuses on the inflation rate when setting interest rates. Second, economic activity indicators were in the focus of the ECB before the financial crisis. Third, over the last decade, the role of economic activity decreased, indicating that inflation is the main driver of monetary policy decisions in the post-crisis period. Fourth, when setting interest rates, central bankers appear to consider more than one model.

*Keywords:* Taylor rule, Bayesian model averaging, model uncertainty, monetary policy, ECB.

*JEL classification:* C11, E43, E58.

---

Chapter 1 has been published as Gross and Zahner (2021). We thank Jesús Crespo-Cuaresma and Florian Huber for raising our interest in the topic of Bayesian model averaging. Furthermore, we are very grateful for the helpful comments by Elisabeth Schulte, Bernhard Herz, Bernd Hayo, Peter Tillmann, Tim Friehe, and the participants of the graduate seminars at the University of Bayreuth and at the Marburg Centre for Institutional Economics (MACIE).

## 1.1 Introduction

> *"The challenge for monetary policy in practice is to retain the virtues of rule-based policy-making, while taking into account the complex, uncertain and constantly evolving environment facing monetary policy-makers."*
>
> — European Central Bank (2001, p. 38)

Due to its pivotal role in monetary policy, the interest rate setting of central banks is one of the most debated topics in the field of macroeconomics. When setting its interest rates, the European Central Bank (ECB) advocates a rule-based approach while retaining room for discretionary interventions. Generally, researchers apply the so-called Taylor rule in order to analyze the interest rate setting (Taylor, 1993). The Taylor rule proposes that inflation and economic activity can well approximate the short-term interest rate. Various studies suggest that this simple rule is a powerful tool to approximate central bank interest rates under special conditions (e.g. Sturm and Wollmershäuser, 2008).

However, there are various reasons why applying a standard Taylor rule yields misleading policy implications. First, it seems far-fetched to restrict decision-makers to a monetary policy rule that captures sufficient information about the real economy in only two variables. Vítor Constâncio, former vice president of the ECB, argues that

> "*the environment in which monetary policy-makers have to act is much more complex than what is assumed in model-based analysis of policy rules. [...] A simple rule that responds to one or two macroeconomic variables and ignores all other indicators of price developments is not able to account for the complexities of the real world.*"
>
> — European Central Bank (2017)

Second, monetary policy decisions are based on incomplete information and uncertainties about the actual state of the economy. Hence, central bankers might analyze a variety of different economic variables and indicators — besides the inflation rate and economic activity — to obtain more accurate information about the state of the economy (Milani, 2008; Lee et al., 2013; Lee et al., 2015).

Third, in reality, central bankers might have more than one model in mind of how the economy functions. The interest rate setting body of the ECB, the Governing Council, consists of the presidents of the euro area national central banks and the Executive Board. These central bankers have different backgrounds that might affect their attitudes towards relevance of certain variables such as inflation and economic activity and the importance of indicators such as bond yields to affect those variables. Although we remain agnostic about the source of heterogeneity, central bankers differ with respect to their social, political, and academic backgrounds. However, aside from potential heterogeneity between the *different* members of the Governing Council, it can further be assumed that heterogeneity exists *within each* central banker, i.e., that more than one reaction function is in her mind. Such heterogeneity leads to different concepts about the transmission of shocks and the interaction of economic agents that might yield deviating policy implications and interest rate recommendations. Therefore, we argue that the standard Taylor rule should be extended to draw more precise inferences for monetary policy.

We contribute to the current monetary policy literature by shedding light on the ECB's monetary policy decisions and analyze the potential shift of priorities due to the global financial crisis in 2008. Our analysis focuses on the following two key factors: Firstly, uncertainty about the form of the central bankers' reaction functions, and secondly, uncertainty about the magnitude of the coefficients included in the specific reaction functions. We base our analysis on real-time data and insights derived from textual analysis of ECB press conference statements.

Employing an empirical Bayesian model averaging (BMA) approach allows us to consider variables besides the ones included in the classical Taylor rule, and to evaluate ∼33,000 model combinations of potential monetary policy determinants. We consider a variety of variables and evaluate all model combinations with respect to the observed data to determine the most likely models. Thereby, we derive the ECB's most likely interest rate determinants. Applying BMA in the context of the ECB's monetary policy is — to the best of our knowledge — a novel approach and addresses a gap in the current literature.

Our key findings are as follows: First, the ECB focuses its decisions mainly on the inflation rate measured by the Harmonized Index of Consumer Prices (HICP). Second, economic activity seems to be a key priority for the central bank before the financial crisis. Third, our results suggest that the importance of economic activity for the ECB's monetary policy decisions decreased over the last decade. Fourth, when setting interest rates, central bankers from the ECB tend to consider more than one model. This finding supports the necessity to use model averaging techniques in order to take model uncertainties in the context of monetary policy into account.

Our paper is structured as follows: In Chapter 1.2, we identify the variables potentially influencing the ECB's monetary policy by a literature review and a textual analysis of ECB communication. Furthermore, we motivate why one should consider model uncertainty in the context of monetary policy. In Chapter 1.3, the BMA approach is discussed. In Chapter 1.4, the data used is discussed. In Chapter 1.5 and Chapter 1.6, the estimation results are explained in detail, and robustness checks are conducted. Chapter 1.7 concludes the paper.

## 1.2 The ECB's Potential Interest Rate Determinants

In 1993, John Taylor proposed that the central bank-set short-term interest rate $i$ can be approximated according to the following rule:

$$i_t = Z_t'\beta + \varepsilon_t, \tag{1.1}$$

where $\beta$ is a vector of reaction coefficients, $Z$ a matrix of macroeconomic variables, and $\varepsilon$ an error term. In its initial specification, Taylor proposed the inclusion of the macroeconomic variables inflation and a measure for economic activity captured in the matrix $Z$.[3] Quite surprisingly, this standard Taylor rule provided a relatively good fit for the Federal Reserve (Fed)'s and ECB's actual interest rate path (Gerdesmeier and Roffia, 2003; Sturm and Wollmershäuser, 2008).

[3] Note that Taylor (1993) assumes a constant real interest rate. However, in most advanced economies there is growing evidence for a decline in real interest rates (Del Negro et al., 2019). We address this decline by dividing the sample period and presume a constant real interest rate within those subsamples.

**Figure 1.1:** Comparison of the actual and Taylor rule-approximated ECB interest rate

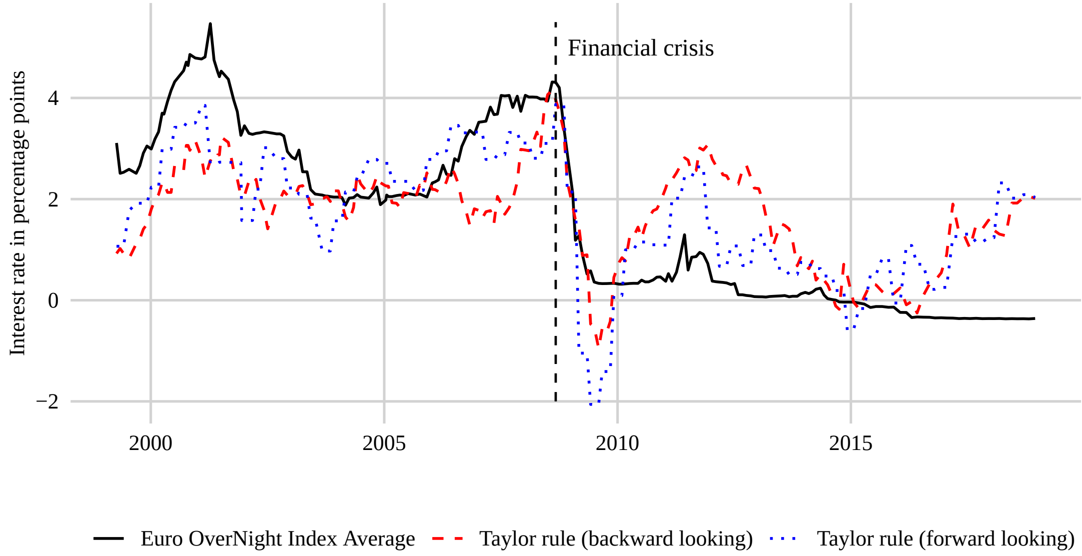


This fit is illustrated in Figure 1.1, where we compare the actual ECB interest rate and its Taylor rule-based approximations for the period from 1999 until 2018.[4] We measure the ECB's interest rate $i$ by the Euro OverNight Index Average (EONIA) rate (black line) and use two Taylor rule-based interest rate approximations as comparisons: one measured using current inflation and output (red dashed line) and one using expected inflation and expected output (blue dotted line).[5] Figure 1.1 provides anecdotal evidence for the following three conclusions. First, both Taylor rules approximate the short-term interest rate (relatively) well until the outbreak of the financial crisis in 2008. Second, after the outbreak, deviations became more severe, suggesting a more restrictive monetary policy stance than actually conducted. The divergence between the actual interest rate and the Taylor rule-based approximations indicates that a single reaction function based on the standard Taylor rule might not be sufficient to approximate the actual short-term interest rate. Therefore, it seems unlikely that two variables are sufficient to properly approximate the short-term interest rate.

Consequently, we argue in line with Milani (2008) that the central bankers in charge

[4] The corresponding regression table can be found in Chapter 1.8.1.

[5] Note that conclusions derived from Taylor rules based on actual data (*backward looking*) or based on expected data (*forward looking*) can vary and, therefore, provide different policy implications (e.g. Svensson, 2003; Gerdesmeier and Roffia, 2003). The implications of forward- and backward-looking specifications are further discussed in Chapter 1.4.

do not seem to follow one single model when setting interest rates. Instead, the ECB's interest rate setting can be best approximated by using a variety of models. Considering different models is equivalent to assume that each central banker has a model about the economy in mind and her model is assigned some probability of being chosen to represent the aggregated Governing Council decision.

We use textual analysis of the ECB's communication as an additional instrument to analyze the ECB's monetary policy reaction functions. To be precise, we analyze introductory statements of the ECB's press conferences, the official communication instrument of the central bank to the general public. Using textual analysis on these statements, we provide evidence to include further variables in the reaction function. A word cloud of the introductory statements (see Chapter 1.8.2) shows that the ECB discusses myriads of different variables and indicators, each, at least implicitly, a potential indication for some reaction by central bankers to the respective variable. Even though the ECB mentions 'inflation' and 'output' in the introductory statements frequently, those two variables do not encompass all of the attention. Also other variables, e.g., related to financial stability and commodity prices, are discussed by the ECB frequently. Therefore, the introductory statements indicate that additional variables might play a role for the ECB.

We do not choose our variables with a view to the Taylor rule, but follow an 'unbiased' approach by using a three-step procedure to identify potential determinants. In a first step, we derive these variables from the monetary policy literature. In a second step, we eliminate variables that were not mentioned in papers analyzing the ECB's monetary policy in a rule-based environment at least once. In a last step, we identify variables that the ECB presidents referred to in their introductory statements. An overview of all potential determinants can be found in Chapter 1.8.3.

Note that two aspects are beyond the scope of this paper: First, we do not examine whether these additional variables enter the ECB's monetary policy framework on their own or as an instrument, i.e., by affecting other variables such as inflation and economic activity. Therefore, the derived variables from the estimated reaction functions do not necessarily reflect the determinants of the central bank's objective function. As a result, causal interpretation should be regarded with caution. Second, we do not claim to estimate the ‘true’ ECB reaction function since the Governing Council’s monetary policy discussions are not provided to the general public. By using exclusively public information, we analyze the public perception of an ECB monetary policy reaction function rather than a private – ECB internal – reaction function.

### 1.2.1 Business Cycle Variables

**Inflation and Output** According to its official mandate, the ECB’s primary objective is to maintain price stability (Article 127 §1 of the Treaty on the Functioning of the European Union). Since 2003, price stability is defined as the medium-term annual growth of the HICP of below but close to two percent (European Central Bank, 2019a). A second, subordinated, objective of the ECB is to support the general economic policies of the European Union (EU) (Article 127 §1 and Article 3 of the Treaty on the Functioning of the European Union). This objective is often interpreted as considering economic activity, i.e., output and unemployment, when setting interest rates. Due to their pivotal roles, it is not surprising that the terms ‘inflation’ and ‘output’ are used frequently in the communication of the ECB. These terms account for more than 1.3% of all words used in press conference statements, i.e., ‘inflation’ is mentioned almost 2000 times and ‘production’ and ‘output’ 142 times. Therefore, we are confident that inflation and output are relevant for the ECB’s reaction functions.

Note that the actual indicators measuring inflation and economic activity are discussed heavily. While various Taylor rules use actual data for inflation and output, Svensson (2003) argues that a standard Taylor rule focusing on backward-looking data is not optimal as those variables affect monetary policy with a lag. Therefore, he suggests to use inflation and output expectations as applied in Sauer and Sturm (2007) and Gerlach (2007). Further, it is ambiguous whether the ECB focuses on HICP inflation, core inflation (Gerlach, 2007), or commodity prices. In our analysis, we use both backward-looking and forward-looking inflation and output as well as core inflation, and commodity prices.

**Unemployment** Article 3 of the Treaty on the Functioning of the EU specifies full employment as one objective of the EU. Therefore, the euro area-wide unemployment rate could be relevant for the ECB's interest rate decision. Furthermore, the state of the labor market can be interpreted as an indicator for economic activity (Molodtsova and Papell, 2012). Textual analysis supports the importance of unemployment. Terms containing the word 'employment' are mentioned more than three times as often (429 times) as the terms 'production' or 'output'.

### 1.2.2 Financial Markets Variables

One of the most prominent extensions of the standard Taylor rule is with respect to the stability of financial markets (Peek et al., 2016). We identified various channels through which financial stability could influence the ECB's interest rate setting (Kaefer, 2014).

**Credit Measures** Credit growth is suspected to impact financial stability via two channels: Firstly, asset bubbles could emerge if credit is massively invested in asset markets. Secondly, credit-financed consumption and investment could lead to a non-sustainable level of debts, which negatively influence economic activity. The term 'credit' is mentioned more than 600 times in the ECB's communication.

**Euro Exchange Rate** If a currency appreciates, exporters become less competitive, driving a decline in output and inflation. This appreciation might lead to capital inflows potentially creating asset bubbles. Holders of the appreciated currency experience financial gains, and consumers adjust inflation expectations. The relevance of the exchange rate is emphasized by the ECB mentioning the exchange rate more than 100 times.

**Euro Area Government Bond Yield** A decrease in the government bond yield eases refinancing cost for countries and is expected to work as a stimulus for the respective economy. Furthermore, Roskelley (2016) stresses the forward-looking component of government bond yields. Terms such as 'government debt', 'government deficit', and 'government bond' are mentioned more than 100 times in the ECB press conferences.

**Stock Prices** The inclusion of asset prices, such as stock prices, follows a similar argumentation as the inclusion of the exchange rate.[6] An increase in asset prices can lead to a rise in output and inflation through increased consumption (consumption smoothing) and investments. Thereby, stock prices include relevant information about future inflation. Terms related to stock and asset prices appeared more than 170 times in ECB communication. Additionally, in the current literature, researchers use market volatility as a measurement for financial stability (Albulescu et al., 2013; Bleich et al., 2013). Since we analyze a wide information set, we include stock prices as well as volatility measures.

**Financial Stress Indicator** The global financial crisis revealed that instabilities in the financial sector can impact financial stability. Hollo et al. (2012) constructed an indicator measuring systemic risk and contemporaneous stability in the financial system, the so-called Composite Index of Systemic Stress (CISS). The index aggregates information beyond the above-mentioned channels about the current instability of the financial system. 'Stress' is mentioned almost 30 times in ECB communication.

**Money Supply** One important component of ECB's policy strategy is the monetary analysis. In the long run, an expansion in the monetary base is expected to drive inflation (Friedman, 1963; Belke and Klose, 2010). Money supply ('M3') is mentioned more than 500 times in the ECB communication analyzed.

### 1.2.3 Further Variables

**Economic Policy Uncertainty** Political and economic uncertainty could encourage individuals to postpone investment and consumption decisions. Aastveit et al. (2013) argue that, in the presence of uncertainty, individuals react more cautiously to interest rate decisions. In order to ensure effective monetary policy measures, central banks would have to act more aggressively to achieve their objectives. Fur-

[6]Note that we do not include property prices as property prices typically adjust slowly due to market frictions and infrequent valuations since transactions are conducted seldom. As a result, the response to central banks' monetary policy measures takes place with a considerable lag. Therefore, property prices may reflect long-term interest rate expectations (e.g., forward guidance) and do not solely react to actual short-term monetary policy changes.

thermore, Philip Lane, the chief economist of the ECB, argues that uncertainty can affect the wage-setting of companies (Lane, 2019). Due to the relevance of uncertainty (mentioned 280 times), we include Baker et al.'s (2016) economic policy uncertainty index for the euro area.

**Trade Deficit** Obstfeld and Rogoff (2005) and Ferrero et al. (2008) argue that trade flows capture information on future currency movements. Hence, the trade balance can be suspected to be a source of inflationary pressures. The ECB mentions 'trade' almost 70 times.

**Interest Rate Smoothing** Note that we do not account for interest rate smoothing as we assume that the central bank monitors a variety of different indicators and variables when setting interest rates. Therefore, the central bank reacts to different economic signals — both positive and negative — that provide mixed interest rate implications indicating a smooth response as positive and negative signals yield an inert interest rate response.[7] Milani (2004) argues that "the explicit introduction of a wider information set is [...] itself a cause of interest rate smoothness".

## 1.3 Estimation Approach

The main objective of this paper is to analyze the ECB's monetary policy under the assumption that central bankers consider a multitude of models containing a wide information set. In this chapter, we briefly summarize the methodology of BMA.[8] Using a BMA approach allows us to account for model uncertainty and to assess and evaluate every feasible model combination that can be constructed from a predefined dataset based on the potentially relevant variables identified in Chapter 1.2.

In previous publications, BMA was already applied to account for model uncertainty in economic contexts, e.g., by analyzing determinants of gross domestic product (GDP) growth (Koop, 2003; Fernandez et al., 2001b), and exchange rate crises

[7]This hypothesis is later supported when analyzing BMA-based monetary policy reaction functions.

[8]We refer the interested reader to more comprehensive BMA literature such as Fernandez et al. (2001a), Milani (2008), Zeugner and Feldkircher (2015), and Moral-Benito (2010). We mainly follow the notation of Moral-Benito.

(Cuaresma and Slacik, 2009). Milani (2008), Lee et al. (2013), and Lee et al. (2015) were the first to use a BMA approach in the context of monetary policy. Lee et al. (2013) estimate a 'meta Taylor rule' — i.e., they characterize monetary policy with reference to the Taylor rule — for the United Kingdom using data from 1972 until 2010 and for Australia from 1972 until 2011. Lee et al. (2015) use a similar approach for deriving monetary policy reaction functions for the United States, analyzing data from 1972 until 2008. Both papers apply the BMA approach in the context of monetary policy to account for uncertainties with respect to the duration of monetary policy regimes and the specification of the Taylor rule. While both do not consider variables beyond the standard Taylor rule, i.e., inflation and economic activity, they also address regime uncertainty — an aspect we do not focus on in our analysis.[9] Whereas Lee et al. (2013) and Lee et al. (2015) estimate a three-digit number of models, we evaluate approximately 33,000 different models and, therefore, consider a wider range of different variables and resulting models. A wide range of variables is also used in Milani (2008) making his approach most similar to ours. In his setup, also variables beyond the 'classical' Taylor rule are included as potential determinants — different to Lee et al. (2013) and Lee et al. (2015), but in line with our approach. However, Milani (2008) applies BMA for the United States considering data from 1979 only until 2001.

In contrast to Milani (2008), Lee et al. (2013), and Lee et al. (2015), we analyze the ECB's monetary policy of the last two decades and take structural changes around the financial crisis and the emergence of the effective lower bound (ELB) into account. Furthermore, we integrate elements from textual analysis to allow for more systematic insights. Our paper is — to the best of our knowledge — the first to use BMA in the context of monetary policy around the financial crisis and, generally, the first to analyze ECB's monetary policy using BMA.

We implement the BMA approach in the following way: First, we specify the ECB's short-term interest rate as the dependent variable $y$. Second, we choose a set of in-

[9] Lee et al. (2013) and Lee et al. (2015) specifically analyze the impact of monetary policy regimes, e.g., related to the stability of the policy responses to economic conditions, policy horizons, or the presidents of the respective central bank. In this paper, we differentiate between the period before the global financial crisis (1999–2008) and the period during and after the crisis (2008–2018). Therefore, we consider two different 'regimes', but do neither address the role of different presidents, nor changes independent of the global financial crisis.

dependent variables potentially influencing the dependent variable ($Z$). The matrix $Z$ consists of all potential regressors that were identified and discussed in Chapter 1.2. Third, we perform the actual BMA estimation and evaluate all possible linear combinations of these regressors.[10] Uncertainty about the (subjective) choice of the 'true model' vanishes as not only one single model is analyzed, but rather all possible model combinations. For every model, the regression coefficients captured in the vector $\hat{\beta}$ are estimated via Bayesian techniques. As a last step, we compute the average coefficient weighted by the respective likelihood of the model. This vector of average coefficients $\hat{\beta}_{BMA}$ can be expressed as

$$\hat{\beta}_{BMA} = \sum_{j=1}^{2^k} \hat{\beta}^j w^j, \tag{1.2}$$

where $\hat{\beta}^j$ is the coefficient estimate of model $j$ and $w^j$ is the respective weight. In the following, $\hat{\beta}$ is referred to as the posterior probability and $w$ as the posterior model probability.

### 1.3.1 Bayesian Model Averaging

The following linear model is considered:

$$y_t = Z_t'\beta + \varepsilon_t, \tag{1.3}$$

$$\varepsilon_t \sim N(0, \sigma^2),$$

where $y_t$ is the interest rate and $Z_t$ is a vector of $k$ explanatory variables at time $t$. $\beta$ is the $k$-dimensional vector of regression coefficients and $\varepsilon_t$ is a vector of error terms, which follows an univariate normal distribution with zero mean and variance $\sigma^2$. With $k$ being the amount of possible regressors, we observe model space $M = \{M_j : j = 1, \ldots, 2^k\}$ and coefficients $\beta = \{\beta^j : j = 1, \ldots, 2^k\}$ with each model $j$ having $k$ individual beta coefficients. By applying Bayes rule, the posterior probability — the distribution of the estimated coefficient vector conditional on one

[10] Note that the integration of interaction terms or other non-linear parameters is still an open research topic in BMA and, therefore, excluded from our analysis.

specific model $j$ and the underlying data — can be derived as follows:

$$p(\beta^j|y, M_j) = \frac{p(y|\beta^j, M_j)p(\beta^j|M_j)}{p(y|M_j)}. \tag{1.4}$$

(1.4) states that the posterior probability $p(\beta^j|y, M_j)$ is calculated by multiplying the likelihood $p(y|\beta^j, M_j)$ with the probability $p(\beta^j|M_j)$ divided by the marginal likelihood $p(y|M_j)$.[11] Note that the marginal likelihood is constant over all models and is, therefore, a multiplicative term. Thus, the marginal likelihood is not in the focus of our analysis since it does not depend on $\beta$, which we seek to examine in this paper.

The inclusion of own information — so-called priors — is one of the key features of Bayesian modeling. In this way, a researcher defines distributional assumptions, e.g., of a coefficient or a model space, before observing the data. Note that the prior choice expresses subjective beliefs and has to be set with caution. In this paper, $p(\beta^j|M_j)$ is referred to as the prior on the parameter space expressing our belief about the probability distribution of $\beta^j$. In Chapter 1.3.2, different model prior specifications are discussed.

As a next step, we aggregate the posterior distributions over the whole model space using an aggregation weight $w^j$. The posterior model probability $p(M_j|y)$ is used as a weight since it indicates the degree of support for model $M_j$. Applying Bayes rule yields this posterior model probability:

$$w^j = p(M_j|y) = \frac{p(y|M_j)p(M_j)}{p(y)} = \frac{p(y|M_j)p(M_j)}{\sum_{s=1}^{2^k} p(y|M_s)p(M_s)}. \tag{1.5}$$

(1.5) states that the posterior model probability — the probability of selecting model $j$ — depends on the marginal likelihood $p(y|M_j)$, the marginal probability $p(M_j)$, and the integrated likelihood $p(y)$. Note that the integrated likelihood does not vary across models and is, therefore, only a multiplicative term. To compute the posterior model probability, a second prior needs to be introduced ($p(M_j)$). This prior specifies the distribution over the model space and expresses our belief about the probability of choosing model $M_j$ before observing the data.

As a next step, a weighted average of all individual posteriors probabilities is com-

[11] Note that $p(y|M_j) = \sum_{s=1}^{2^k} p(y|\beta^s, M_s)p(\beta^s|M_s)$.

puted to obtain one full posterior distribution. As a weight, the posterior model probability from (1.5) is used. Hence, the full posterior distribution can be expressed as follows:

$$p(\beta_{BMA}|y) = \sum_{j=1}^{2^k} p(\beta^j|y, M_j)p(M_j|y). \tag{1.6}$$

This full posterior distribution allows us to analyze coefficients across all models. We examine economic relevance of the included variables by estimating the expected value $E(\beta_{BMA}|y)$ and the variance $V(\beta_{BMA}|y)$ of each coefficient. Both moments can be derived from (1.6) as follows (Moral-Benito, 2010; Koop, 2003):

$$E(\beta_{BMA}|y) = \sum_{j=1}^{2^k} E(\beta^j|y, M_j)p(M_j|y). \tag{1.7}$$

$$V(\beta_{BMA}|y) = \sum_{j=1}^{2^k} V(\beta^j|y, M_j)p(M_j|y) + [E(\beta^j|y, M_j) - E(\beta^j|y)]^2 p(M_j|y). \tag{1.8}$$

To identify relevant variables, we analyze the posterior inclusion probability (PIP) and rank the most relevant regressors based on their fit. The PIP for variable $h$ can be computed as follow:

$$PIP_h = p(\beta_h \neq 0|y) = \sum_{\beta_h \neq 0} p(M_j|y). \tag{1.9}$$

A variable with a high PIP indicates that the variable is included in a variety of relevant models and can, thus, be considered robust. We define variables with a PIP $> 0.15$ as robust.

### 1.3.2 Definition of Priors

In this chapter, we specify the priors on the parameter space $p(\beta^j|M_j)$ from (1.4) and on the model space $p(M_j)$ from (1.5).

**Prior on the Parameter Space** We follow Koop (2003) and specify a normal distribution with zero mean for the distribution of the coefficient to put as less (subjective) information on the distribution as possible before observing the data. For the variance, we apply the $g$-prior proposed in Zellner (1986), who introduced an additional parameter $g$ into the variance structure.[12] The vector of the estimated coefficients $\beta$ of model $j$, therefore, follows the following normal distribution:

$$\beta^j|\sigma^2, M_j, g, X \sim N(0, \sigma^2 g(X_j' X_j)^{-1}). \quad (1.10)$$

In line with the BMA literature, the standard deviation parameter $\sigma$ is assumed to be equal in all models and is set as an uninformative prior, as proposed in Fernandez et al. (2001a) and Zeugner and Feldkircher (2015)[13] :

$$p(\sigma) \propto \frac{1}{\sigma}. \quad (1.11)$$

Therefore, the expected value of $\beta^j$ can be expressed as follows (Zeugner and Feldkircher, 2015):

$$E(\beta^j|M_j, g) = \frac{g}{1+g}\hat{\beta}^j_{OLS}. \quad (1.12)$$

Equation (1.12) shows that by using the $g$-prior, the expected value of the coefficients can be expressed as a convex combination of the ordinary least squares (OLS) estimator $\hat{\beta}^j_{OLS}$ and the prior mean (zero). By specifying $g$, the researcher indicates how much importance she puts on the prior belief. A small $g$ expresses a high weight of the prior mean. In the case $g \to 0$, the expected value of the coefficient converges to the prior mean (zero). In the case $g \to \infty$, the estimated value approaches the OLS estimator, neglecting the prior completely. This prior structure yields a likelihood $p(y|M_j, X, g)$ that is similar to $R^2$ and includes a penalty for large models.
In the literature, three different possibilities for the specification of $g$ are considered primarily, namely the unit information prior (UIP) proposed by Kass and Raftery (1995) ($g = t$), the risk inflation criterion (RIC) by Foster and George (1994) ($g = k^2$) and the benchmark risk inflation criterion (BRIC) by Fernandez et al. (2001a) ($g = max(t, k^2)$). In our estimation, we apply the BRIC prior as it combines the UIP and the RIC by using the maximum $g$ of both candidates. We

[12]The popularity of the g-prior variance specification is mainly due to the facts that (1) a closed-form solution for the posterior distributions exists reducing computational issues, (2) the variance of the coefficient only depends on the scaling parameter $g$ as $\sigma$ is equal in all models (Moral-Benito, 2010), and (3) a penalty for large models is included.

[13]Note that the prior choice p($\sigma$) does not influence the estimation results since $\sigma$ is equal in all models and, therefore, has the same implications for every model (Koop, 2003).

thereby put as little importance on the prior as possible. In Chapter 1.6.1, we show that our results are robust to different prior specifications.

**Prior on the Model Space** The prior on the model space $p(M_j)$ defines the expectation about the number of regressors the researcher believes to be included in the true model. For example, if the prior is set to five the researcher expects the dependent variable to be most accurately explained by five independent variables. Hence, the researcher can use the prior on the model space to express a preference for smaller or larger models. We assume that the model size $\Xi$ follows a binomial distribution, specified as $\Xi \sim Bin(k, \theta)$, where $\theta$ is the prior inclusion probability for each variable. Therefore, model $M_j$ with $k$ regressors has a prior model probability of

$$p(M_j) = \theta^{k_j}(1-\theta)^{k-k_j}. \tag{1.13}$$

In the following, we consider two approaches for implementing this prior on the model space, namely through a binomial distribution and a binomial-beta distribution. Both priors have the advantage of being easy to implement but have the disadvantage of neglecting multicollinearity issues, i.e., that the probability that a regressor will be included in a model is observed separately.[14] In other words, the inclusion or elimination of one variable does not change the probability of any other regressors being included. Our approach to address the multicollinearity issue is discussed in Chapter 1.5.

**(a) Binomial (Uniform) Model Prior.** Using a binomial approach, the expected model size can be expressed as:

$$E(\Xi) \equiv m = k\theta. \tag{1.14}$$

Milani (2008) sets $\theta = 1/2$, allocating the highest probability (and the expected value) to models that contain $k/2$ variables. An alternative is to set $\theta < 1/2$. This specification increases the likelihood of smaller models and could take limitations

[14] Multicollinearity can lead to outcomes that highly correlated variables are included in the models causing biased estimation results.

in human cognitive abilities into account. In other words, it can be suspected that central bankers tend to consider primarily small models. We implement several specifications for $\theta$ accounting for a lower probability of large models. Note that our results indicate a tendency towards small models, independent of the specification of $\theta$.

**(b) Binomial-beta Prior.** Ley and Steel (2009) suggest the use of a hyperprior on the inclusion probability $\theta$. This specification makes $\theta$ random, in contrast to the binomial distribution, where $\theta$ is fixed. Ley and Steel suggest to use a beta distribution for the hyperprior, i.e., $\theta \sim Beta(a, b)$. The use of such a binomial-beta prior leads to an expected model size of

$$E(\Xi) \equiv m = \frac{a}{a+b}k. \tag{1.15}$$

Ley and Steel (2009) propose $a = 1$ and $b = (k - m)/m$ expecting the researcher to specify — similar to the binomial prior — only the expected model size $m$. Choosing $m = k/2$ leads to $a = b = 1$, which yields the following (flat) model size probability distribution:

$$p(M_j) = \frac{1}{k}. \tag{1.16}$$

(1.16) indicates that the selected binomial-beta distribution results in a posterior probability that is equal for each model size. Thereby, it reduces the subjective influence concerning the expected model size which minimizes the impact of the prior choice by the researcher. While the probability distribution is centered around the expected model size for a binomial distribution, in the case of a binomial-beta distribution, the probability distribution of the posterior is flat for all models.

## 1.4 Data

In this chapter, we discuss the data used in our BMA analysis. Note that we focus on real-time data, which allows a more comprehensive understanding of the perceived central bankers' reaction to macroeconomic data. Orphanides (2001) provides evidence that monetary policy implications based on revised (not real-time) data are inaccurate as the data used by policy-makers and researchers does not align. He argues that revisions constitute new information about previous data points, not available to policy-makers at that point in time. Since the ECB acknowledges that macroeconomic variables were subject to significant revisions (European Central Bank, 2010), it is essential to use real-time information. We obtain such real-time data from the ECB's own Real-Time Database (RTD).

Due to the following three factors, we argue that the RTD represents an accurate information set available to ECB central bankers. First, using the RTD allows to merge low frequency macroeconomic data such as GDP or inflation with higher granular financial data such as asset prices or exchange rates. Second, it seems reasonable to assume that policy-makers primarily consider own information. Third, the RTD contains the latest macroeconomic data, relevant for the central bankers, as revisions are usually published on the day before the monetary policy decisions. However, the RTD does not contain data for all the variables discussed in the previous chapters. In particular, expectation data of macroeconomic variables, such as inflation, economic activity, and unemployment, are not included. We obtain expectation data from the ECB's Survey of Professional Forecasters (SPF). The remaining data is obtained from various sources (see Chapter 1.8.3).[15]

The interest rate set by the ECB's Governing Council is the main refinancing operations (MRO) rate. However, the presence of the ELB and the ECB's unconventional monetary policy measures do not favor the MRO as an appropriate representation of the short-term interest rate. Instead, we use two alternatives: First, we consider the EONIA rate. The EONIA rate is applied in Taylor rules, e.g., in Fendel and Frenkel (2006) and Castro (2011). Second, we consider the Wu and Xia shadow rate for the euro area (Wu and Xia, 2016) since it captures additional information from unconventional monetary policy measures such as forward guidance or asset purchase programs.

[15]For most data, the only transformation was to compute annual growth rates to reduce stationarity. One exception is the output gap. We calculated the output gap by applying the Hodrick-Prescott (HP) filter on real-time data for real GDP. Note that this is different from the method applied in Taylor (1993). Instead of real-time data, he uses an end-of-sample measure and constructs his data of the output gap from the deviation of actual real GDP from a constant trend real GDP.

**Figure 1.2:** Comparison of the shadow rate by Wu and Xia (2016) and the EONIA rate

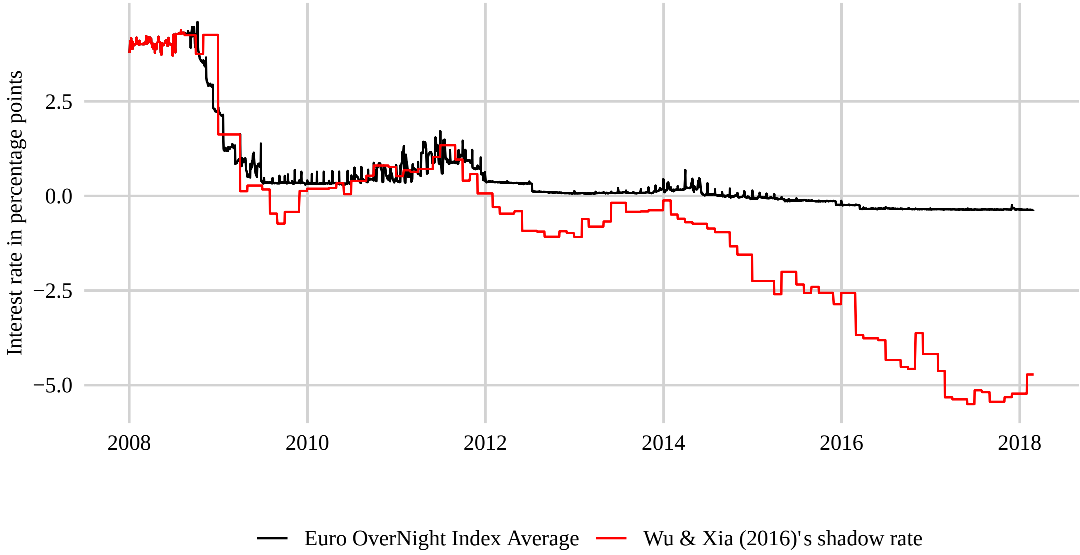


Figure 1.2 highlights only minor deviations between the EONIA and the shadow rate from 2008 until 2011. However, the divergence between the two rates increased from 2011 onwards — with the EONIA being restricted by the ELB — and reached a gap of more than five percentage points from 2017 onwards. The shadow rate ($s$) is applied as the dependent variable for the period after the outbreak of the global financial crisis when unconventional monetary policy measures were introduced. We set the date for the outbreak of the financial crisis to September 15th, 2008. On that day, the bankruptcy of Lehman Brothers was officially declared. Therefore, the short-term interest rate $y$ is comprised of the EONIA rate $i$ and the shadow rate $s$, i.e.:

$$y_t = \begin{bmatrix} i_t \\ s_t \end{bmatrix} \begin{matrix} \text{if } t < \text{September 2008,} \\ \text{if } t \geq \text{September 2008.} \end{matrix} \tag{1.17}$$

## 1.5 Results

In the following chapter, we discuss the main results of our BMA analysis. We apply the BRIC $g$-prior on the parameter space and a flat binomial-beta distribution on the model space. These specifications are chosen in order to put as little (subjective) information on the priors as possible. As described in Chapter 1.4, the EONIA rate and the Wu and Xia shadow rate have been combined as the dependent variable. As we discuss in the next chapter, our main findings are robust across different prior specifications. We run the regressions in R, using the BMS package by Zeugner and Feldkircher (2015). To account for multicollinearity, we use Cuaresma and Slacik's (2009) approach and neglect all regressors with a correlation[16] higher than $|0.6|$.[17] In total, we estimated $2^k = 2^{15} \approx 33{,}000$ different models for the time period from April 1999 until March 2018 aggregating the respective model-specific estimation results to obtain one average effect for every regressor.

**1999–2018** The main results of the regressions are shown in Table 1.1. The first column shows the regressor, the second the PIP — the aggregated probability of the models including the respective variable —, and the third the post mean — the average marginal effect — with the standard deviations denoted in brackets.

**Table 1.1:** Regression results

| | *1999 – 2018* | | *1999 – 2008* | | *2008 – 2018* | |
|---|---|---|---|---|---|---|
| | PIP | Post Mean | PIP | Post Mean | PIP | Post Mean |
| HICP Inflation Rate | 1.000 | 1.095<br>(0.169) | 0.736 | 0.313<br>(0.216) | 0.941 | 0.776<br>(0.273) |
| Unemployment (expected) | 1.000 | -0.666<br>(0.111) | 0.239 | -0.068<br>(0.136) | | |
| Output Gap (expected) | | | 0.986 | 1.708<br>(0.438) | | |
| Output Gap (actual) | | | 0.967 | 0.500<br>(0.179) | | |
| Observations | | 216 | | 114 | | 102 |

*Note:* Only robust variables with a PIP $\geq 0.15$ are presented.

The main results for the period from 1999 until 2018 are the following: First, the HICP inflation rate is included in all relevant models, indicated by a PIP of 100%. Therefore, our results suggest that the current inflation rate is indeed primarily considered when the ECB sets interest rates. The central bank reacts to a 1%

[16] We recognize the limitation of using bivariate correlations as a representation for multivariate relationships. For further discussions, see Hayo (2018).

[17] Actual unemployment, expected inflation, and money growth were excluded from the regression due to strong correlation. We have conducted a range of robustness checks to see if our choice has altered the outcomes. The primary findings appear to be selection independent.

increase in the annual inflation rate by increasing the interest rate on average by 1.1%. However, the inflation coefficient's standard error equals 0.17. Thus, the Taylor principle cannot be confirmed with reasonable statistical significance. Second, expected unemployment has a PIP of 100% as well. The coefficient is negative, indicating that the central bank reacts to an increase in the expected unemployment rate with expansionary monetary policy. Third, no further variables are considered in the majority of the models, i.e., no variable has a PIP > 15%.

**Figure 1.3:** Ten top models: Whole period

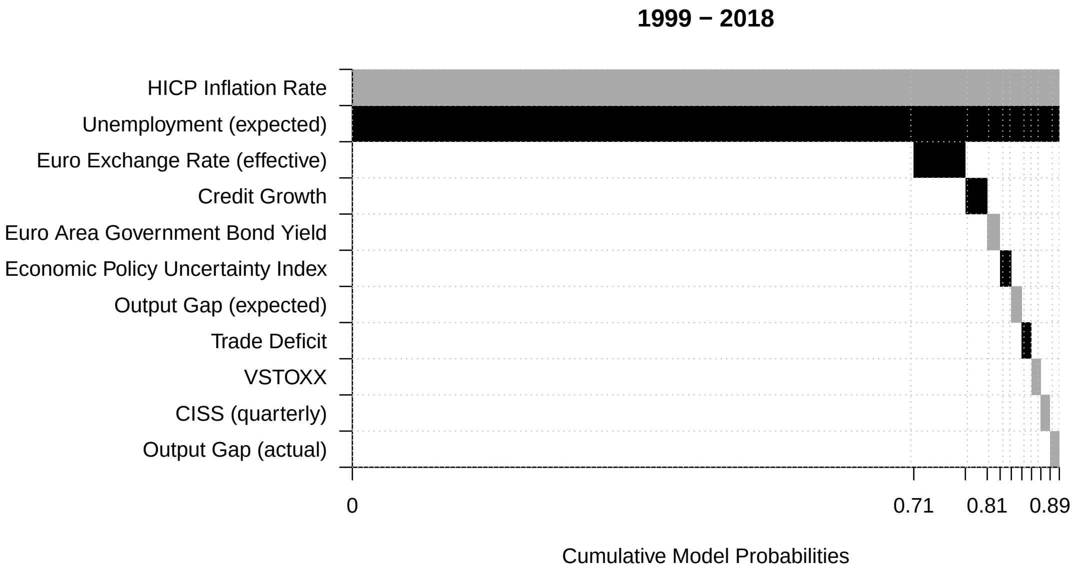


Next, we consider individual models. Figure 1.3 provides an overview of the ten 'best-performing' models ranked by their respective model inclusion probabilities. The higher the model inclusion probability, the higher the likelihood that this specific model represents the 'true' model. The cumulative model probabilities — an aggregated likelihood of models considered — are shown on the horizontal axis and the relevant regressors on the vertical axis, i.e., the variables that are included in the models with the highest likelihood. Gray color indicates a positive sign of the respective coefficient, while black color indicates a negative sign. The model that can explain the data in the most precise manner (the 'top model') has a model probability of 0.71, i.e., the best individual model out of $2^{15}$ models has a likelihood of 71%. This specific model includes only the two (robust) variables, namely inflation and expected unemployment with the expected signs. The model with the second highest likelihood additionally includes the effective euro exchange rate that yields a likelihood of approximately 7%. The next best model with the inflation rate, unemployment, and credit growth has a model likelihood of approximately 3%. Thus, the figure indicates that model probabilities are not evenly distributed across a large number of heterogeneous models. Restricting the analysis to one single model — a common approach in classical model-selected Taylor rules — would neglect relevant model probabilities. Therefore, using our ten top models instead of a two variable Taylor rule increases the cumulative inclusion probability from 71% to almost 90%.

**Figure 1.4:** Frequency of ECB's determinant communication

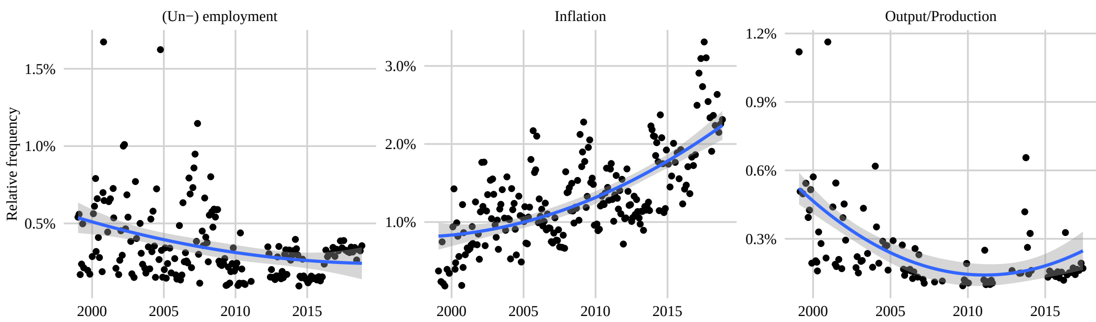


**1999–2008** A crucial question is whether the ECB altered its monetary policy strategy after the financial crisis. One approach to detect a potential systemic change in the perceived ECB reaction function is to observe the communication of the central bank. Figure 1.4 illustrates the relative frequency of terms related to inflation and economic activity mentioned in ECB press conferences over time. Analyzing the relative frequency, we find evidence for a gradual shift from unemployment and output towards inflation. To account for this potential shift, we separate the time periods during and after the crisis explicitly from the period before the financial crisis.[18] The main results for the pre-crisis period are displayed in Table 1.1 in columns four and five: First, the output gap seems to be the main determinant of the ECB's interest rate. Both the expected output gap — based on expectations — and the actual output gap — based on previous data — are robust and significant. Both

[18]Note that we use the term 'post-crisis' for the time period after the financial crisis although it actually includes both the period of the crisis itself as well as the post-crisis period.

variables are included in almost all models with the expected sign: An increase in the output gap leads, on average, to an increase in the interest rate. The coefficient for the expected output gap (1.7) is higher than the coefficient for the actual output gap (0.5). This reaction is in line with the ECB's official objective to support economic activity in the euro area. Quite interestingly, the coefficient for the actual output gap equals Taylor's initial calibration from 1993. However, Taylor includes only one measure of economic activity — a key difference to our empirical results. Note that the focus on the output gap is in line with the communication of the ECB. As discussed previously, the fraction of output-related terms decreases over time (see Figure 1.4).

Second, the HICP inflation rate is again robust (PIP = 0.74), although with a lower coefficient. Note that for the pre-crisis period the Taylor principle can be significantly rejected. Third, the expected unemployment rate is included in a minority of the models (PIP = 0.24) with the expected sign. Unemployment is only included in two out of the ten top models (see Figure 1.8). Therefore, our results suggest that some central bankers — prior to 2008 — saw unemployment as an important determinant to consider in the context of monetary policy but the majority favored the output gap as a measure for economic activity. Note that, overall, the unemployment rate coefficient is not statistically significant. Fourth, further indicators, such as the exchange rate or stock market prices, only enter the minority of the models (PIP $\leq$ 0.15) and are, thus, not considered robust.

To summarize, our results suggest that the main determinants of the ECB's monetary policy for the period before the financial crisis are related to the business cycle. These findings are in line with the official mandate of the ECB to account for both the development of inflation and economic activity. However, the Treaty on the Functioning of the European Union specifies that the two objectives are hierarchical in a sense that the ECB mainly focuses on inflation and only subordinately on the business cycle. Our result do not confirm this hierarchy. We find a stronger focus on economic activity than on inflation, indicated by a higher posterior mean, a higher PIP, and that the Taylor principle cannot be confirmed.

Focusing on individual models (see Figure 1.8), the top model includes the (actual and expected) output gap and the HICP inflation rate (PIP = 31%). The ten top models yield a cumulative model probability of 0.62. Therefore, the results provide again evidence to favor averaging techniques when analyzing the monetary policy of the ECB.

**2008–2018** Next, we analyze the ECB's monetary policy in a post-crisis context. The main results for the post-crisis period are displayed in Table 1.1 in columns six and seven. The findings suggest a shift in the ECB's monetary policy strategy. On the one hand, inflation stays highly robust. In the post-crisis period, the PIP is even higher and indicates that HICP inflation is included in almost all models. The sign of the coefficient is again positive and the magnitude of the coefficient is higher compared to the pre-crisis period. Therefore, our results suggest that the ECB shifts its focus towards the inflation rate after the bankruptcy of Lehman Brothers. For this period, we can neither reject nor confirm the Taylor principle. On the other hand, none of the economic activity measures — actual and expected output gap, and unemployment — are robust for the post-crisis period (PIP $\leq 0.15$).
The results suggest that the focus of the ECB shifted from considering both inflation and economic activity in the pre-crisis period to solely considering the HICP inflation rate in the post-crisis period. This finding is supported by evidence from the ECB's communication, showing that, in the post-crisis period, the ECB mentions terms related to inflation more often and terms related to economic activity less often. The ten top models have a cumulative model probability of 80%. The top model only includes HICP inflation and yields a likelihood of 0.55 (see Figure 1.9).

**Short Summary** To sum up all periods, our results cannot confirm that the ECB does not account for the inflation rate. Both in normal times and in times of crisis, inflation is included in the majority of the relevant models. Furthermore, we find evidence that the ECB reacts to the expected unemployment rate and the output gap. To be precise, our results suggest that the ECB seems to focus its monetary policy decisions after Lehman bankruptcy mainly on the HICP inflation rate, while before the financial crisis both inflation and economic activity measures seem to be relevant. However, the size of the inflation coefficients provide evidence that the Taylor principle might not be fulfilled.

**Figure 1.5:** Model sizes of the different time horizons

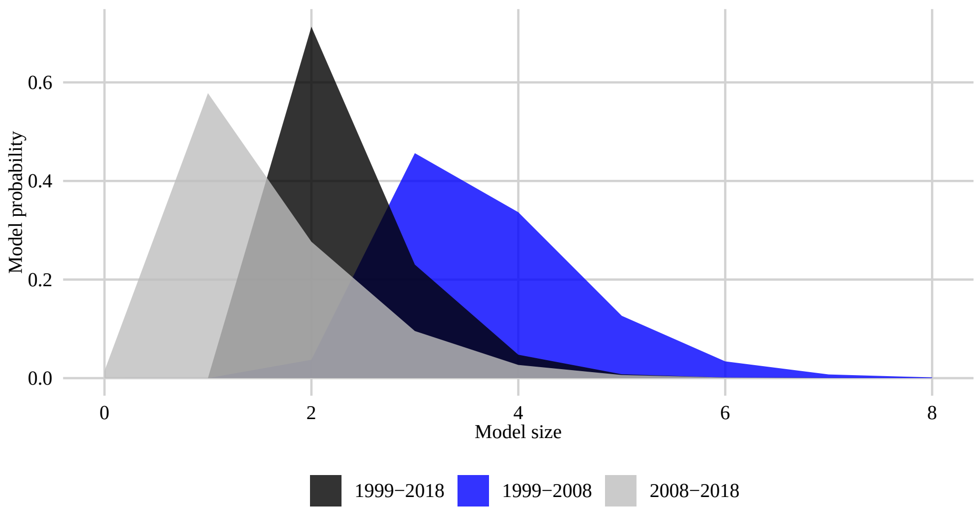


**Model Size** Next, we draw inferences on the number of included regressors in the interest rate setting by evaluating the distributions of the posterior model sizes. Figure 1.5 plots the posterior model sizes for the different time periods. The black density plot refers to the period from 1999 until 2018 (whole period), the blue one from 1999 until 2008 (pre-crisis), and the gray one from 2008 until 2018 (post-crisis). The figure indicates that the distributions of the model sizes vary slightly with respect to their means. Depending on the specification, the average number of determinants in the monetary policy reaction functions are 2.4 (whole period), 3.7 (pre-crisis), and 1.6 (post-crisis). Hence, these results correspond to our previous findings in Table 1.1, i.e., that the ECB seems to consider fewer variables after the financial crisis when setting interest rates. However, independently of the time period considered, large models with more than five regressors seem unlikely. Hence, the ECB's monetary policy is best approximated using medium-sized models with between one and five variables.

**Table 1.2:** Model fit evaluation

| | MSE | AIC | BIC |
|---|---|---|---|
| Taylor rule (backward looking) | 5.30 | 376.72 | 390.31 |
| Taylor rule (forward looking) | 4.95 | 361.38 | 374.97 |
| BMA 1999 – 2018 | 3.55 | 302.16 | 339.54 |
| BMA 1999 – 2008 & 2008 – 2018 | 1.89 | 163.16 | 200.54 |

**Model Fit** Next, we analyze the model fit of the ten top models via an in sample prediction. The main results are presented in Table 1.2 and can be summarized as follows: First, an improvement in explanatory power for the whole period can be reached by using BMA. If we compare the BMA-derived ten top models with the two standard Taylor rules (see Figure 1.1), the goodness of fit improves according to the mean squared error (MSE) between 28% and 33%. This improvement could potentially be caused by overfitting. However, it is unlikely that overfitting is the sole driver of the better fit since we demonstrated that the derived ten top models include only a few variables. To further examine overfitting, we evaluate our results using the Akaike information criterion (AIC) and the Bayesian information criterion (BIC) as alternative measures for goodness of fit. For both criteria, we find similar improvements.[19] Therefore, we argue that the improvement in fit is predominantly driven by increased information provided by the additional regressors.

**Figure 1.6:** Comparison of BMA-derived approximation vs. actual interest rate

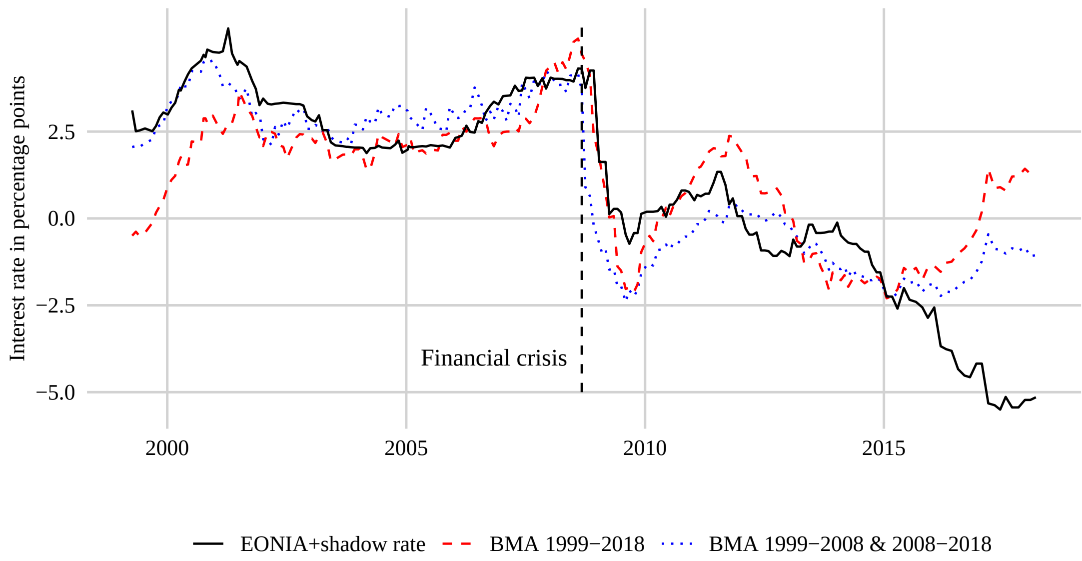


Second, the separation between pre- and post-crisis periods leads to a further substantial improvement in fit as indicated in Table 1.2, with a decrease in the MSE by more than 60%. Moreover, we find similar improvements when controlling for

[19] AIC and BIC are more appropriate criteria compared to the MSE. AIC and BIC introduce a trade-off between overfitting and underfitting by including a penalty increasing with the number of regressors. To analyze the robustness of our results, we conservatively estimate both measurements with the maximum number of possible variables instead of the actual maximum model size.

the higher amount of variables in both the AIC and BIC. Figure 1.6 compares the BMA-derived interest rate approximations with the actual short-term interest rate captured by the EONIA and the shadow rate. Until 2015, both BMA approximations provide similar results with little variation. Afterwards, the two BMA approximations diverge substantially. The increase in the BMA-derived interest rate for the whole period (red dashed line) suggests more restrictive monetary policy measures in that period. The specification separating the whole period in the pre- and post-crisis period (blue dotted line) supports the expansionary monetary policy stance to a higher degree. The separated specification advocates an interest rate higher than the shadow rate, but still below the ELB. At the end of our observation period, the gap between our two BMA approximations is more than 2%. Note that inflation as the main determinant in the post-crisis period has steadily increased since the end of 2015 leading to a divergence from the shadow rate.

## 1.6 Robustness Checks

In this chapter, we discuss the robustness of our results. Robustness checks are conducted with respect to (1) prior modifications, (2) a varying date of the beginning of the financial crisis, and (3) the use of a different dependent variable. We conclude that our results are robust across a multitude of model specifications.

### 1.6.1 Priors

In a first robustness check, we modify the priors. As mentioned in Chapter 1.3.1, priors are chosen subjectively by the researcher. Therefore, it is essential to conduct robustness checks to different prior specifications.

**Prior on the Model Space** In the benchmark case, we apply a flat binomial-beta model prior. However, as described previously, one could assume that central bankers favor small models. Therefore, as a robustness check, we apply a different binomial-beta prior, thereby putting a higher weight on small models ($m = 2$). Results are illustrated in Chapter 1.8.5. The findings are similar to the benchmark case. If the whole time period is considered, inflation and the expected unemployment rate are again the only robust variables and the respective posterior means are almost identical. When separating the pre-crisis and post-crisis period, the findings with respect to the PIPs and coefficient estimates do not vary substantially, as well. To summarize, if a higher weight is put on smaller model sizes, the key findings remain the same.

**Prior on the Parameter Space** Next, we alter the prior on the parameter space. In the benchmark case, we apply the BRIC g-prior. Note that this prior combines the RIC and the UIP. As long as $k^2 > N$, the BRIC equals the RIC. In our case ($k = 15$ and $N = 216$), the RIC dominates the UIP. Hence, applying the RIC rather than the BRIC does not alter the outcomes. The results for a robustness check using the UIP prior are shown in Chapter 1.8.5. The findings are (again) almost identical with respect to the PIP and the posterior means compared to the benchmark case. To summarize, applying different priors on the parameter space leads to similar results.

### 1.6.2 Starting Date of the Crisis

In our benchmark case, the beginning of the financial crisis is set to the day of Lehman's bankruptcy, where also macroeconomic data indicated an upcoming recession. However, one can also date the beginning of the financial crisis 13 months earlier, precisely on August 9, 2007. On this day, the inter-banking market in the euro area broke down and the ECB provided additional funding for banks (Stark, 2010). We, therefore, perform another robustness check altering the starting date of the financial crisis. Results are shown in Chapter 1.8.5. In the pre-crisis period, the robust variables are almost identical to the benchmark case and most of the PIPs and post means are similar, with two exceptions: Firstly, the expected unemployment rate is not robust anymore. Secondly, the inflation coefficient decreases and becomes insignificant. For the post-crisis period, the PIP of inflation increases from 0.94 to 1.00 and the coefficient increases almost twofold to 1.22. To summarize, redefining the starting date of the crisis pronounces our main result for the post-crisis period even more, namely that inflation becomes the most relevant determinant in the post-crisis period as the PIP and the coefficient size increase.

### 1.6.3 Endogenous Variable

In this robustness check, we substitute Wu and Xia's (2016) shadow rate with the EONIA rate in the post-crisis period, i.e., we use the EONIA rate for the whole period from 1999–2018 as our endogenous variable. The findings of the robustness check are shown in Chapter 1.8.5. For the whole period, the results are qualitatively similar. However, the coefficients are — due to the lower variance in the endogenous variable — smaller. Note that the results for the pre-crisis period remain the same. For the post-crisis period, the findings from the benchmark case seem to be confirmed qualitatively, as inflation still appears to be a main determinant of the ECB's monetary policy. Besides the inflation rate, the expected unemployment rate is robust and significant. Few other variables tend to be robust — the output gap, commodity prices and core inflation —, but are neither economically relevant nor statistically significant. In summary, our benchmark results are robust and qualitatively independent of the selection of the endogenous variable.

## 1.7 Conclusion

Over the last decade, the standard Taylor rule, using inflation and economic activity to approximate the ECB's short-term interest rate, has lost substantial explanatory power. This divergence might indicate that central bankers consider other variables beyond those suggested in Taylor (1993) and employ different models when setting interest rates. In this paper, we mainly attribute this divergence to model uncertainty. We analyze a wide array of potential determinants that we derive from the literature and textual analysis of the ECB's press conference statements. Using a BMA approach enables us to assess and evaluate every feasible model combination constructed from the variables identified. This approach has — to the best of our knowledge — not been applied previously in the context of the ECB's monetary policy and, therefore, addresses a gap in the current literature. Our derived reaction functions aim to provide clarity for understanding the ECB's monetary policy and allow both researchers and the general public to draw conclusions about potential determinants of the ECB's monetary policy. By separating pre-crisis and post-crisis periods, we account for a potential shift in the central bank's monetary policy strategy.

Our key findings are the following: First, our results indicate that, irrespective of the period analyzed, the inflation rate is a robust determinant. However, our analysis suggests that the inflation coefficient is not statistically significantly different from one. Therefore, we do not find evidence that the Taylor principle is fulfilled. Second, the importance of inflation in terms of robustness and coefficient magnitude increased over time. In fact, for the last decade, inflation appears to be the only robust determinant. This result seems to be in accordance with the communication of the ECB. Third, we find that the robustness of the output gap has decreased over the last decade. Fourth, small single-digit models, including between two and seven variables, approximate the ECB's monetary policy most precisely. Fifth, the distribution of model probabilities shows that no single model can sufficiently explain the observed data. This finding reaffirms using model averaging methods when evaluating monetary policy of central banks rather than selecting one single — e.g., standard Taylor rule-based — model. Nonetheless, we can explain most of the variation in the interest rate by analyzing only ten models.

This paper provides a first analysis of applying model averaging techniques for the ECB interest rate setting. Future research could extend the model averaging approach as follows: First, dilution priors could be incorporated into BMA applications as an alternative approach to account for multicollinearity issues. Second, whereas BMA has already been applied to monetary policy of other central banks (e.g. Lee et al., 2015), our approach incorporates novel techniques — such as textual analysis and analyzing a broader range of variables — and enables the determination and comparison of monetary policy reaction functions of different central banks. Applying this approach to the interest rate setting of other central banks, such as the Fed or the Bank of England, could provide an interesting comparison of similarities and differences of monetary policy strategies of global central banks. Third, future research could put a higher focus on accounting for heterogeneity, i.e., by considering macroeconomic developments of specific euro area countries and not euro area aggregates. Finally, in our analysis, we neglect regime uncertainty in the spirit of Lee et al. (2015). Accounting for different regimes may add another dimension to our understanding of how central banks conduct monetary policy.

## 1.8 Appendix

### 1.8.1 Taylor Rule Estimate

**Table 1.3:** Regression results: Taylor rule

| | *Dependent variable* | |
|---|---|---|
| | EONIA | |
| | (1) | (2) |
| Inflation | 0.975*** | |
| | (0.102) | |
| GDP | 0.287* | |
| | (0.170) | |
| Inflation (expected) | | 2.862*** |
| | | (0.254) |
| GDP (expected) | | 0.774*** |
| | | (0.080) |
| Constant | −0.010 | −4.248*** |
| | (0.202) | (0.427) |
| Observations | 224 | 221 |
| $R^2$ | 0.339 | 0.560 |
| Adjusted $R^2$ | 0.333 | 0.556 |
| *Note:* | *p<0.1; **p<0.05; ***p<0.01 | |

### 1.8.2 Textual Analysis of ECB Communication

**Figure 1.7:** Wordcloud of ECB communication (bigrams)

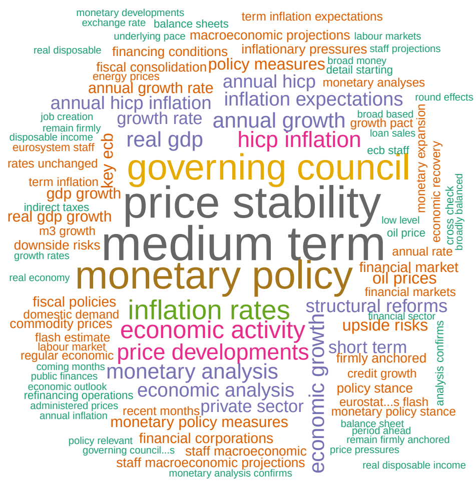


Source: ECB press conference introductory statements.

### 1.8.3 Summary of Variables

| Name | Source | Transformation |
|---|---|---|
| Dependent Variables | | |
| EONIA Rate | Datastream | - |
| Euro Area Shadow Rate (Wu & Xia (2016)) | Quandl.com | - |
| Business Cycle | | |
| HICP Inflation Rate | ECB (RTD) | Annual growth |
| Exp. Inflation Rate | ECB (SPF) | Annual growth |
| Output Gap | ECB (RTD) | HP transformed real GDP |
| Exp. Output Gap | ECB (SPF) | HP transformed exp. GDP |
| Unemployment | ECB (RTD) | - |
| Exp. Unemployment | ECB (SPF) | - |
| Core Inflation Rate | ECB (RTD) | Annual growth |
| Commodity Prices | IMF | Monthly growth |
| Financial Markets | | |
| Effective Euro Exchange Rate | ECB (RTD) | Monthly growth |
| Stock Prices (Euro Stoxx 50) | Datastream | Monthly growth |
| Stock Market Volatility (VSTOXX) | Datastream | Monthly growth |
| Credit Volume | ECB | Monthly growth |
| Euro Area Government Bond Yield | FRED | Monthly growth |
| CISS | ECB | Weekly + monthly growth |
| Money Supply (M3) | ECB (RTD) | Annual growth |
| Further Indicators | | |
| Uncertainty Index | policyuncertainty.com | Monthly growth |
| Trade Deficit | ECB (RTD) | Quarterly growth |

### 1.8.4 Ten Top Models (Pre- and Post-Crisis)

**Figure 1.8:** Top models: Pre crisis

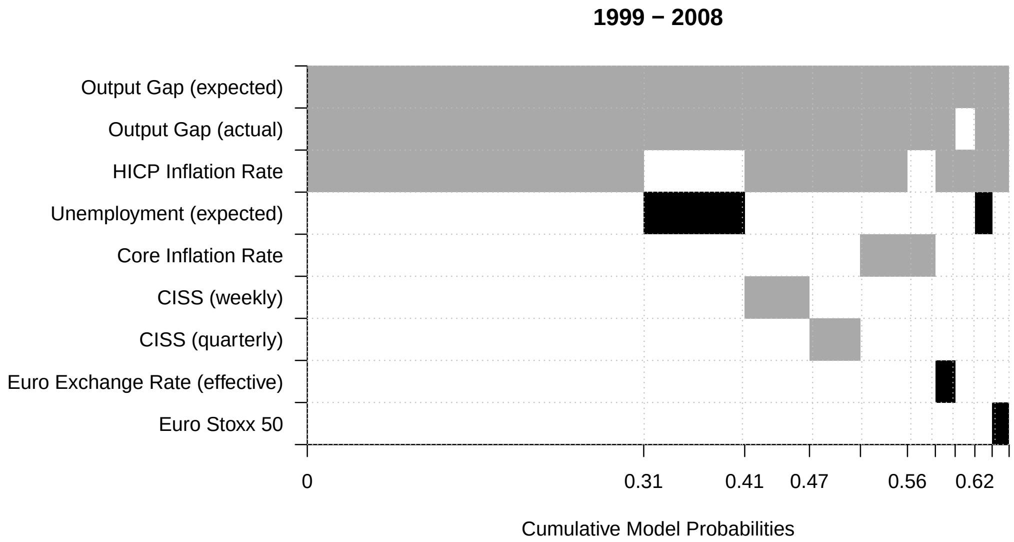


**Figure 1.9:** Top models: Post crisis

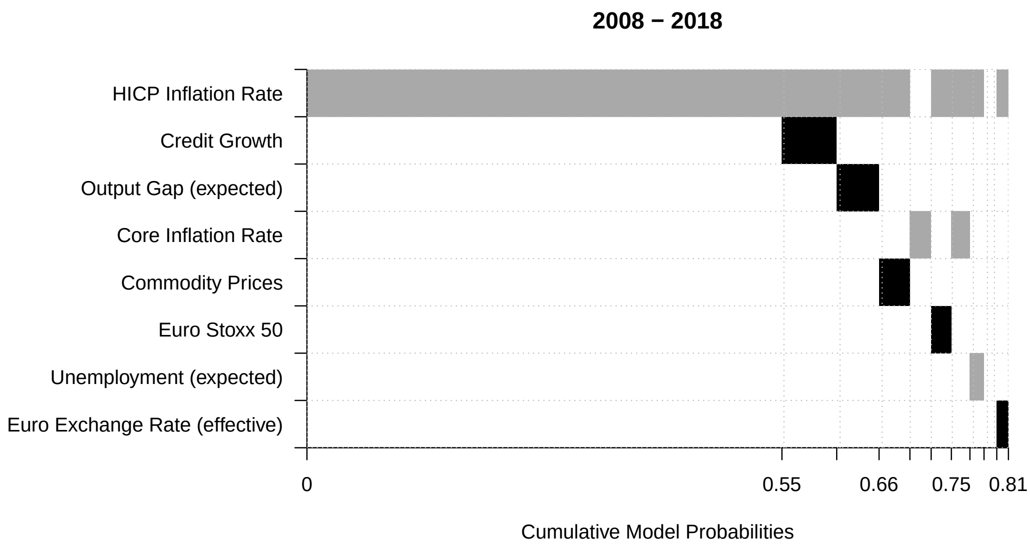

### 1.8.5 Robustness Checks

Robustness: Prior on the model space

| | *1999 – 2018* | | *1999 – 2008* | | *2008 – 2018* | |
|---|---|---|---|---|---|---|
| | PIP | Post Mean | PIP | Post Mean | PIP | Post Mean |
| HICP Inflation Rate | 1.000 | 1.100<br>(0.169) | 0.719 | 0.313<br>(0.221) | 0.931 | 0.770<br>(0.281) |
| Unemployment (expected) | 1.000 | -0.666<br>(0.110) | 0.228 | -0.069<br>(0.140) | | |
| Output Gap (expected) | | | 0.980 | 1.712<br>(0.456) | | |
| Output Gap (actual) | | | 0.952 | 0.490<br>(0.190) | | |
| Observations | | 216 | | 114 | | 102 |

*Note:* Only robust variables with a PIP $\geq 0.15$ are presented.

Robustness: Prior on the parameter space

| | *1999 – 2018* | | *1999 – 2008* | | *2008 – 2018* | |
|---|---|---|---|---|---|---|
| | PIP | Post Mean | PIP | Post Mean | PIP | Post Mean |
| HICP Inflation Rate | 1.000 | 1.094<br>(0.169) | 0.748 | 0.310<br>(0.211) | 0.948 | 0.778<br>(0.267) |
| Unemployment (expected) | 1.000 | -0.666<br>(0.111) | 0.254 | -0.068<br>(0.113) | | |
| Output Gap (expected) | | | 0.989 | 1.700<br>(0.427) | | |
| Output Gap (actual) | | | 0.977 | 0.502<br>(0.172) | | |
| Observations | | 216 | | 114 | | 102 |

*Note:* Only robust variables with a PIP $\geq 0.15$ are presented.

Robustness: Starting date of the crisis

| | *1999 – 2018* | | *1999 – 2007* | | *2007 – 2018* | |
|---|---|---|---|---|---|---|
| | PIP | Post Mean | PIP | Post Mean | PIP | Post Mean |
| HICP Inflation Rate | 1.000 | 1.095<br>(0.169) | 0.224 | 0.091<br>(0.189) | 1.000 | 1.216<br>(0.219) |
| Unemployment (expected) | 1.000 | -0.666<br>(0.111) | | | | |
| Output Gap (expected) | | | 0.971 | 1.890<br>(0.568) | | |
| Output Gap (actual) | | | 0.874 | 0.474<br>(0.349) | | |
| Observations | | 216 | | 101 | | 115 |

*Note:* Only robust variables with a PIP $\geq 0.15$ are presented.

Robustness: EONIA

| | *1999-2018* | | *1999-2008* | | *2008-2018* | |
|---|---|---|---|---|---|---|
| | PIP | Post Mean | PIP | Post Mean | PIP | Post Mean |
| Unemployment (expected) | 1.000 | -0.694<br>(0.059) | 0.238 | -0.067<br>(0.135) | 0.990 | -0.191<br>(0.048) |
| HICP Inflation Rate | 0.998 | 0.440<br>(0.094) | 0.736 | 0.313<br>(0.216) | 0.856 | 0.200<br>(0.101) |
| Output Gap (expected) | | | 0.985 | 1.708<br>(0.438) | | |
| Output Gap (actual) | | | 0.967 | 0.498<br>(0.179) | 0.609 | -0.143<br>0.130 |
| Commodity Prices | | | | | 0.999 | -0.061<br>(0.010) |
| Core Inflation Rate | | | | | 0.299 | -0.164<br>(0.281) |
| Observations | | 216 | | 114 | | 102 |

*Note:* Only robust variables with a PIP $\geq 0.15$ are presented.

# Chapter 2

# A Model for Central Bank Digital Currencies: Implications for Bank Funding and Monetary Policy

**Abstract**

We develop a dynamic stochastic general equilibrium model to study the impact of central bank digital currencies (CBDCs) on the financial sector. We focus on the effects of interest- and non-interest-bearing CBDCs during financial crises, also on the effective lower bound. In addition, we analyze the role of central bank funding and a rule-based flexible interest rate on CBDC. We find that, in times of crises, CBDCs can crowd out bank deposits and negatively affect bank funding. However, this crowding-out effect can be mitigated if the central bank chooses to provide additional central bank funds or to disincentivize large-scale CBDC accumulation through low CBDC interest rates.

*Keywords:* CBDC, financial sector, monetary policy, disintermediation, DSGE.

*JEL classification:* D53, E42, E58, G21.

---

Chapter 2 is a discussion paper. An earlier version of this work has been published as Gross and Schiller (2021). We are very grateful for the helpful comments by Katrin Assenmacher, Alexander Bechtel, Lea Bitter, Carl-Andreas Claussen, Sonja Davidovic, Peter Dittus, Bernhard Herz, Stefan Hohberger, Dirk Niepelt, Werner Röger, David Tercero-Lucas, Mauricio Ulate, Taojun Xie, Johannes Zahner, Corinne Zellweger-Gutknecht, the participants of the MSY seminar at the ECB, the participants of the CBDC Research Forum at the University of Bern, the participants of the Annual Meeting 2021 of the Verein für Socialpolitik, and the participants of the 1st Annual Central Bank Conference on Development Economics in the Middle East and North Africa.

## 2.1 Introduction

The advent of Bitcoin and other private monies, such as global stablecoins, have raised concerns among central banks worldwide. If such cryptocurrencies gain significant market shares, monetary policy transmission and monetary sovereignty could be impaired (European Central Bank, 2020a). In addition, the use of cash as a means of payment — the only form of central bank money available for citizens — is currently declining. Consequently, dependence on private sector payment infrastructures is increasing. In particular in advanced economies, central banks consider issuing retail central bank digital currencies (CBDCs) — that is digital central bank money for private agents — to guarantee payment resilience in an increasingly digital environment, avoid private sector natural monopolies in the payment market, and strengthen monetary sovereignty in the face of new competitors, such as global stablecoins and foreign CBDCs (European Central Bank, 2020a; Brainard, 2021). To a certain extent, a retail CBDC can be considered a substitute for cash. However, unlike cash, CBDC presumably imposes no storage cost, can be transferred comfortably (e.g., via mobile phones), and is less likely to be stolen or lost.

Despite the apparent potential of CBDC, central bankers remain cautious. They fear that a CBDC could threaten financial stability by facilitating (digital) bank runs and disintermediating the financial sector. In this context, disintermediation is defined as a customer-induced substantial conversion of bank deposits into CBDC. As commercial banks rely on deposits to fund their lending business, deposit outflows increase their funding costs and lead, *ceteris paribus*, to a decline in loan volume, investment, and overall economic activity. While, in general, the academic literature on the effects of CBDCs on the financial sector is growing remarkably, more research on their impact on bank funding is needed, particularly (i) on the effects of different CBDC remuneration and (ii) on the role of central bank refinancing. Further, (iii) the monetary policy implications of CBDCs remain underresearched. From a central bank perspective, CBDCs can provide an additional monetary policy tool that can increase monetary policy efficiency by allowing for negative rates and, in the absence of cash, circumvent the effective lower bound (ELB). Currently, there are no simulations of different CBDC remuneration designs or analyses of their impact on the ELB of nominal interest rates.

In this paper, we address these two gaps by developing a New Keynesian dynamic stochastic general equilibrium (DSGE) model with a specific focus on CBDC and the financial sector. In contrast to existing models, our model accounts for the inherent risk of bank deposits during times of financial crises and includes (different degrees of) central bank refinancing for banks. We use this model to assess CBDC-specific dynamics and transmission effects during a financial crisis. This paper (i) studies the options for the central bank to combat potential disintermediation of the financial sector and (ii) analyzes the effects of using a CBDC as a policy instrument. In particular, we consider two different forms of CBDCs — an interest-bearing CBDC and a non-interest-bearing CBDC — with different implications for the ELB.

We build on the model proposed by Gertler and Karadi (2011), a framework that consists of a financial sector, a public sector, different types of producers, and homogeneous households. In their cashless model, bank funding solely consists of households' deposits and accounts for a moral hazard problem. This rigidity increases the persistence of financial shocks, that is, it introduces a financial accelerator effect that mimics the shock persistence of the global financial crisis.

We expand their model such that our framework exhibits necessary features for analyzing CBDC. First, to allow for active portfolio decisions, households no longer automatically provide their deposits to banks based on the moral hazard constraint but instead based on their utility maximization. We introduce heterogeneity in the forms of savings in terms of liquidity, remuneration, and risk, and assume that households choose their savings portfolio based on these differences. We explicitly account for the risk of bank deposits by introducing a discount factor on the expected return on bank deposits, which decreases with the level of debt in the financial sector and the profits of banks. The intuition behind this modeling approach is that households perceive bank deposits as risky when financial sector debt is high and profits are low. They fear that banks could become bankrupt and, thus, in the absence of a deposit insurance scheme, their deposits could be lost. Second, to capture the central bank's prominent role in bank funding and account for additional central bank policies, we introduce the option of central bank loans for commercial banks. These loans are similarly constrained by the bank's moral hazard problem, thus, keeping the financial accelerator effect intact. Third, we introduce a CBDC, which can be remunerated, as an additional option for households' portfolio decisions. We assume that, in terms of liquidity, it is a perfect substitute for bank deposits, but, as central bank money, it exhibits no counterparty risk.

We calibrate the models with and without CBDC such that their steady states are identical and focus our analysis on the resulting dynamics — that is we deliberately abstract from potential steady state effects of a CBDC introduction. Our calibration of conventional parameters closely follows Gertler and Karadi (2011) with two exceptions, namely related to government expenditures and the interest rate on bonds, which are both calibrated based on euro area data. The additional parameters introduced specifically in our model are mainly calibrated to match data on bank funding.

We show that, given the assumption that during a financial crisis bank deposits are perceived as risky, the presence of a CBDC substantially reduces bank funding and, thus, increases the disintermediation of the financial sector. To secure bank funding, the central bank can compensate losses in deposits by providing additional central bank funds. Assuming full allotment, a CBDC does not impair bank funding, but only affects its composition. Consequently, for both interest- and non-interest-bearing CBDCs, the central bank can stabilize the financial sector and mitigate CBDC-specific disturbances in the real economy. If an interest-bearing CBDC can circumvent the ELB, we find substantial macroeconomic improvements for the entire economy. However, these improvements are not directly linked to a CBDC and changes in households' saving behavior. Instead, due to potentially negative interest rates, the increased room for monetary policy mitigates disturbances after a crisis. Relaxing the assumption of full allotment, the resulting imperfect replacement of deposits with funds from the central bank opens up a channel for CBDC to the real economy. Then, the disintermediation of commercial banks negatively impacts investment, the build-up of capital, and production. In this case, a CBDC indeed destabilizes the financial sector and negatively affects the entire economy. Using the remuneration on CBDC as a policy tool, the central bank can mitigate adverse effects for the financial sector and the real economy by disincentivizing substantial CBDC accumulation. A negative remuneration on CBDC, for example, could render

CBDC less attractive compared to its alternatives, thereby reducing the demand for CBDC.

Our paper contributes to the growing literature on CBDCs and their impact on the financial sector. For studying these effects, Bindseil (2020) provides a starting point. In his paper, he uses a balance sheet exercise to define CBDC-specific channels that could affect the financial sector. First model-based analyses study such potential adverse effects in greater detail and analyze the interlinkages of a CBDC with the financial sector. Keister and Sanches (2019) use a new monetarist model with centralized and decentralized markets to conclude that a CBDC might increase banks' funding costs and crowd out deposits. Fernández-Villaverde et al. (2020b) analyze CBDCs in a Diamond and Dybvig (1983)-type model and find that the central bank faces a CBDC trilemma where a socially efficient solution, price stability, and financial stability cannot be achieved simultaneously. Brunnermeier and Niepelt (2019) provide a generic model with money and liquidity and show that — given certain assumptions — a CBDC introduction only alters the composition of bank funding and not its total size. Also using a Diamond and Dybvig (1983)-type model, Fernández-Villaverde et al. (2020a) find that a CBDC does not alter the equilibrium allocation of bank funding. However, in times of crises, the central bank becomes a deposit monopolist potentially endangering maturity transformation. Chiu et al. (2019) also study a model with centralized and decentralized markets and find that a CBDC improves efficiencies in the financial sector, as banks lose market power. In an extreme scenario, a CBDC can then even lead to an increase in banks' lending activities. Andolfatto (2021) uses an overlapping generations model with monopolistic banks and finds that a CBDC might reduce banks' monopoly profits but does not necessarily lead to disintermediation of the financial sector. CBDCs might even increase financial stability, as deposits could expand due to higher deposit interest rates. Barrdear and Kumhof (2021) build a monetary-financial DSGE model and study the steady state effects of an interest-bearing CBDC. Even if the transition would lead to a crowding out of bank deposits, they find that production could increase significantly.

We contribute to this literature on financial sector implications of CBDCs in the following manner. First, we provide a micro-founded model to study the potential adverse effects on bank funding in times of financial crises when deposits are perceived as risky. Second, we analyze implications for the financial sector based on different CBDC remuneration designs.

Our paper also relates to the literature on the implications of CBDC for monetary policy. Dyson and Hodgson (2016) and Bindseil (2020), amongst others, argue that a CBDC can provide substantial monetary stimulus during a severe recession, as, in the absence of cash, CBDC interest rates can overcome the ELB and feature negative rates. Mancini-Griffoli et al. (2018) discuss how CBDCs impact the transmission channels of monetary policy measures and obtain different conclusions. To study transmission channels in detail and in the absence of empirical data, first model-based approaches have been used. Meaning et al. (2021) use a stylized model and conclude that monetary policy transmission would not change substantially, but, for a given change in policy instruments, the efficiency of the transmission might increase. Analyzing the transmission with their DSGE model, Barrdear and Kumhof (2021) find that a CBDC would improve the central bank's ability to stabilize the business cycle. Ferrari et al. (2020) examine monetary transmission in an open economy DSGE model. They conclude that a CBDC increases the size of international spillover shocks and that a national CBDC can decrease monetary policy autonomy in foreign economies.

We contribute to extant literature by studying and comparing the effects of interest-bearing and non-interest-bearing CBDC designs, with a particular focus on their implication for the ELB on nominal interest rates. Further, we highlight the role of interest rate spreads and the allotment of central bank money as monetary policy tools to mitigate CBDC-specific destabilizing effects.

Our results are important for at least three reasons. First, our model simulation provides valuable insights for the ongoing discussions on how to design a CBDC to prevent destabilizing effects for the financial sector. If the central bank is willing to provide a substantial amount of additional central bank loans to commercial banks, CBDC-induced losses in bank funding can be offset. This policy eliminates the need for restrictive designs, such as upper limits on CBDC holdings, as proposed by Panetta (2018). Further, we show that designing a CBDC with a flexible and potentially negative interest rate provides central banks with an effective tool to govern the demand for CBDC. This tool can be used, amongst others, to prevent CBDC-specific disintermediation of the financial sector during times of financial distress. Second, in the absence of empirical data, our model-based analysis sheds light on the general economic impact of a CBDC. We highlight the transmission of financial shocks with CBDCs. Our model provides a microfounded framework to study the potential disintermediation of the financial sector. By accounting for the perceived risk of bank deposits in times of crises, we observe a liquidity effect — that is, households substitute bank deposits with CBDC for liquidity purposes. Third, the results of our CBDC simulation are relevant for central bankers, who perceive CBDCs as an additional instrument for their monetary policy toolkit. In particular, the European Central Bank (ECB) considers a CBDC introduction also for monetary policy reasons (European Central Bank, 2020a)). Our simulations of interest- and non-interest-bearing CBDCs and, in particular, our focus on the ELB provide a starting point to adequately compare the monetary policy implications of different CBDC remuneration designs.

The remainder of the paper is structured as follows. Chapter 2.2 discusses our model. Chapter 2.3 explains and motivates the model calibration. Chapter 2.4 analyzes alternative versions of the model with non-interest-bearing CBDC (2.4.1), with interest-bearing CBDC (2.4.2), with and without full allotment (2.4.3), and with different interest rate rules on CBDC (2.4.4). Chapter 2.5 concludes.

## 2.2 Model

Our model builds on the closed economy New Keynesian framework by Gertler and Karadi (2011). We substantially rework the utility maximization of households, financial intermediaries' funding, and the role of the central bank. In this chapter, we focus on a detailed discussion of our adaptions.[20] The basic structure of our model is depicted in Figure 2.1.

[20]For an in-depth presentation of the other model parts, we refer to Gertler and Karadi (2011) and for a detailed comparison of the models to Chapter 2.6.2.

**Figure 2.1:** Model structure

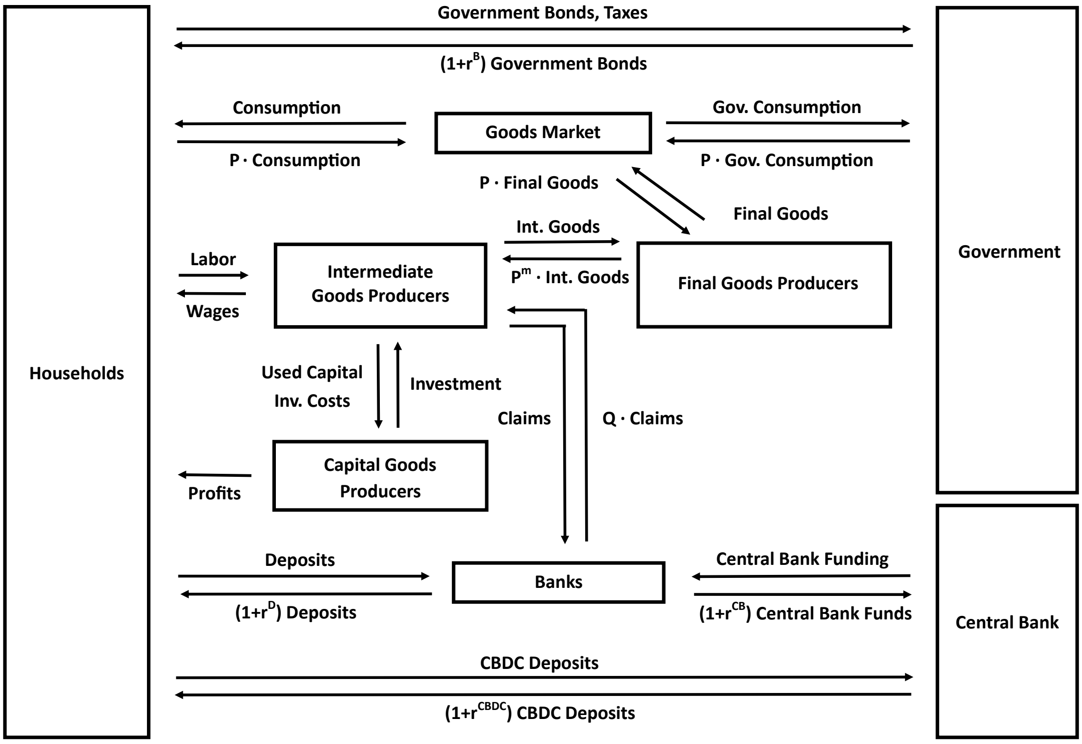


Banks obtain funds from households and the central bank and act exclusively as intermediaries, thereby providing funds for intermediate goods producers. Following Gertler and Karadi (2011), we assume that banks can default and divert obtained funds. The consequent moral hazard that arises places an endogenous limit on banks' balance sheets and restricts their ability to collect funds. While Gertler and Karadi (2011) determine the amount of deposits solely based on banks' economic performance, we determine the amount of bank deposits by households' optimal portfolio choice. We assume that households perceive commercial bank money as risky, particularly in times of financial distress. Households have an incentive to substitute bank deposits with less risky alternatives. They acquire government bonds and CBDC that, additionally, differ in terms of liquidity and remuneration. Further, note that we assume a cashless society.

Intermediate goods producers use intermediated funds to buy capital goods from capital goods producers who face investment adjustment costs. Production requires labor and capital. Competitive monopolistic final goods producers buy intermediate goods, repackage them, and sell them on the goods market to either households or the government.

### 2.2.1 Households

There is a continuum of identical and infinitely lived households that supply labor ($L$), consume goods ($C$), and save for consumption in the next period. They save either via CBDC ($CBDC$), deposits ($D$), or government bonds ($B$). They do not invest in the production sector due to their lack of expertise. We assume that households choose their portfolio in each period without any adjustment costs and not based on love of variety. Instead, the three forms of saving differ in terms of the three dimensions remuneration, liquidity, and risk (see Table 2.1).

**Table 2.1:** Comparison of bank deposits, CBDC, and government bonds

| | **Remuneration** | **Liquidity** | **Risk** |
|---|---|---|---|
| **Bank deposits** | Intermediate | Means of payment | Risky |
| **CBDC** | Low | Means of payment | Riskless |
| **Government bonds** | High | No means of payment | Riskless |

First, with regard to remuneration, deposits pay the real interest rate $r^D$, CBDC pays $r^{CBDC}$, and bonds pay $r^B$ with $r^B \geq r^D \geq r^{CBDC}$.[21] Second, with regard to liquidity, CBDC and bank deposits are perfect substitutes. As both can be used as a means of payment, they generate utility by providing liquidity services. We assume that government bonds do not provide liquidity services, as liquidation is costly and takes time and government bonds are not a means of payment. Third, with regard to risk, CBDC and government bonds are perceived as riskless and bank deposits as risky.

The households' (aggregate) maximization problem can be written in the following manner:

$$\text{(2.1)} \qquad \max\, E_t \sum_{i=0}^{\infty} \beta^i \Big( ln(C_{t+i} - hC_{t+i-1}) + \frac{\Upsilon}{1+\Gamma} \, (D_{t+i} + CBDC_{t+i})^{1+\Gamma} - \frac{\chi}{1+\phi} L_{t+i}^{1+\phi} \Big),$$

where $\Upsilon$ and $\chi$ denote the relative utility weights of real money balances ($CBDC$

[21]In our model, we use this interest rate relation to match data before the outbreak of the global financial crisis and the initiation of substantial asset purchase programs that pushed government bond yields close to, and partially even below, zero.

and $D$) and labor, respectively; $\Gamma$ is the elasticity of money balances, $\phi$ the Frisch elasticity of labor supply, $h$ the habit parameter for consumption, and $\beta$ the intertemporal discount factor. Note that we use a money-in-the-utility-function specification (Sidrauski, 1967; Rotemberg, 1982).[22]

Households believe that banks could go bankrupt and, then, their deposits would be lost. The probability for this event is $1 - \psi$. Note that we abstract from deposit insurance schemes in our analysis.[23] The expected payout of bank deposits can be expressed as

$$(1 - \psi_t)0 + \psi_t(1 + r_t^D)D_t = \psi_t(1 + r_t^D)D_t. \tag{2.2}$$

Hence, the risk can also be expressed as a discount factor on bank deposits. Thus, households' (aggregate) budget constraint can be written in the following manner:

$$\begin{aligned} C_t + D_t + CBDC_t + B_t = w_t L_t + \Pi_t &+ (1 + r_{t-1}^D)\psi_{t-1}D_{t-1} \\ &+ (1 + r_{t-1}^{CBDC})CBDC_{t-1} + (1 + r_{t-1}^B)B_{t-1}, \end{aligned} \tag{2.3}$$

where $w$ is the real wage rate and $\Pi$ income from the ownership of both non-financial (capital goods producers) and financial firms (banks) net of lump-sum taxes $T$. The resulting first-order conditions are derived in Chapter 2.6.1.

The discount factor $\psi$ is increasing in the amount of bank deposits ($D$) and additionally depends on the level of stress in the financial sector, as indicated by losses in banks' equity ($N$):

$$\psi_t = 1 - \left(\frac{D_t}{F_t^*}\right)^{\Omega_D} - \frac{\bar{N} - N_t}{\bar{N}}\Omega_N. \tag{2.4}$$

Banks receive external refinancing both from households and the central bank. $F^*$

[22] Alternatives to our specification would be a cash-in-advance or a shopping-time specification. Apart from slight differences caused by the cross product of consumption and liquidity, these alternatives can be formally equivalent (Feenstra, 1986). We choose this approach to account for the observed large-scale accumulation of money that cannot be justified by precautionary liquidity holdings for future consumption.

[23] Today, deposit insurance schemes are set up to address the risk of commercial bank money and to avoid that, in the case of bankruptcy of a commercial bank, depositors face substantial losses. However, deposit insurance schemes are not available in all countries, and commercial bank money is only secured until a specific threshold. Future research could analyze the interaction of deposit insurance schemes with CBDCs.

denotes the maximum volume of external refinancing implied by the moral hazard in the financial sector (see Chapter 2.2.2). $D/F^*$ is the share of deposits in external refinancing. $\Omega_D$ denotes the elasticity of $\psi$ to changes in bank deposits, while $\Omega_N$ is a scaling parameter and defines the impact of changes in banks' equity $N$.

**Figure 2.2:** Relationship between bank deposits and the discount factor

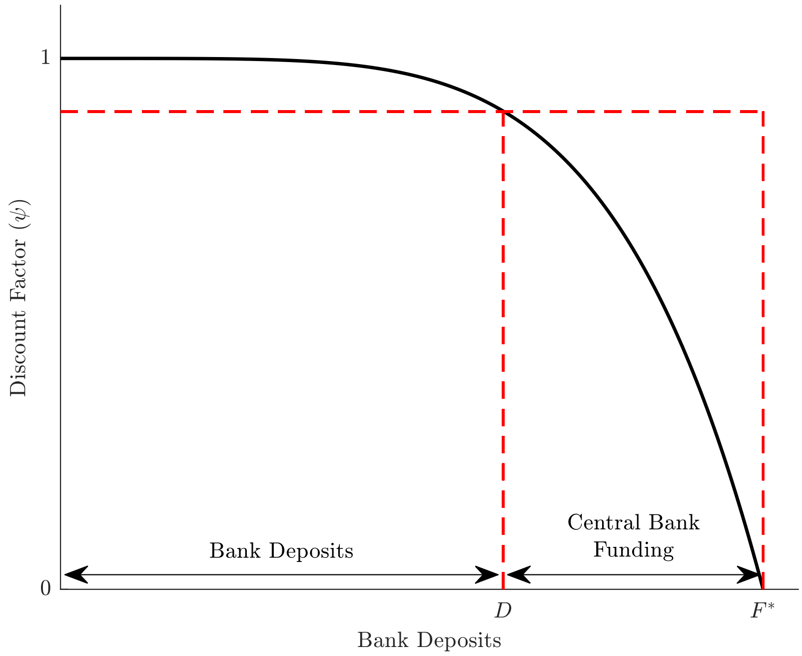


As depicted in Figure 2.2, there is a negative relationship between bank deposits ($D$) and the discount factor $\psi$. When $D$ approaches the maximum amount of external refinancing ($F^*$), where households fear a diversion of their deposits (see Chapter 2.2.2), they perceive deposits as more risky and the discount factor drops. When $\psi$ decreases, such that the expected utility from holding deposits is lower relative to alternative assets, households seek less risky alternatives. In other words, a reduction in $\psi$ can be interpreted as a reduction in the remuneration of bank deposits; subsequently, households decrease their bank deposits. The reduction in $D$ induces banks to demand additional central bank funds in order to secure their lending activities.[24]

[24]Note that we assume that banks always receive the maximum funding ($F^*$). Therefore, if bank deposits decline, a commercial bank demands and receives additional funds from the central bank. This assumption also implies that banks always own sufficient collateral to provide in exchange for additional central bank funds. We relax this assumption in Chapter 2.4.3.

Households perceive this more prominent role of the central bank as a stabilizing factor that lowers the risk in the financial sector. $\psi$ rises up to the point at which households are indifferent between commercial bank money and its alternatives, taking into account the three dimensions remuneration, liquidity, and risk. The elasticity $\Omega_D$ impacts the illustrated curve by shifting it to or away from the upper right corner. Higher values for $\Omega_D$ allow for a higher share $D/F^*$ that households tolerate before they perceive bank deposits as risky. Thus, the calibration of $\Omega_D$, impacts the composition of banks' external refinancing. We use this parameter to calibrate steady-state deposits and central bank funding according to empirical data (for details, see Chapter 2.3).

In addition, $\psi$ depends on the term $\Omega_N \cdot (\bar{N} - N)/\bar{N}$. Thus, we assume that a reduction of banks' equity below its steady state $\bar{N}$ signals financial stress to households and lowers households' trust in commercial banks and, therefore, the discount factor. We use this term to scale the initial impact of the simulated financial crisis on deposits.

### 2.2.2 Banks

Banks use their equity, households' deposits, and funds received from the central bank to acquire claims on intermediate goods producers. The expected return on their investment $r^K$ depends on the performance of intermediate goods producers and is realized by a transfer of any revenues or losses in the next period. Banks pay back households' deposits and central bank funds with the *ex-ante* known nominal interest rates $i^D$ and $i^{CB}$.

Banker $j$ accumulates wealth $N_j$. Wealth can be interpreted, as the banker's equity, while deposits and central bank funds $R_j^{CB}$ represent external refinancing $F_j$. Therefore, banker $j$'s balance sheet relation is given by:

$$Q_t S_{jt} = N_{jt} + D_{jt} + R_{jt}^{CB} = N_{jt} + F_{jt}, \tag{2.5}$$

where $S_j$ captures $j$'s financial claims, priced $Q$, against the production sector. Banker $j'$s equity depends on interest expenses and interest income:

$$N_{jt+1} = (1 + r_{t+1}^K)N_{jt} + (r_{t+1}^K - r_t^D)D_{jt} + (r_{t+1}^K - r_t^{CB})R_{jt}^{CB}. \tag{2.6}$$

Note that a banker's equity is driven by the interest rate spreads — the premia $r^K_{t+1} - r^D_t$ and $r^K_{t+1} - r^{CB}_t$. Banker $j$ intermediates funds as long as the premia are non-negative, which results in the two following participation constraints:

$$E_t \beta \Lambda_{t,t+1}(r^K_{t+1} - r^D_t) \geq 0, \tag{2.7}$$

$$E_t \beta \Lambda_{t,t+1}(r^K_{t+1} - r^{CB}_t) \geq 0, \tag{2.8}$$

where $\beta\Lambda_{t,t+1}$ is the discount factor derived from the first-order conditions of households (see Chapter 2.6.1), as we assume that bankers are part of the household sector, following Gertler and Karadi (2011). In this framework, households consist of a constant fraction of bankers and workers. Each banker might change profession with a worker in each period with a certain probability, thereby transferring all earnings to the household. Households send out new bankers and equip them with start-up funds. This exit-and-entry-mechanism ensures that, in the absence of shocks, the aggregate equity of all bankers does not increase. These assumptions ensure that bankers cannot solely satisfy the demand for funds by intermediate goods producers with their equity and render external refinancing redundant (Gertler and Karadi, 2011). Banker $j$ maximizes the expected terminal wealth, $V_j$, given by

$$V_{jt} = E_t \sum_{i=0}^{\infty} (1-\theta)\theta^i \beta^{i+1} \Lambda_{t,t+i+1}(N_{jt+i+1}), \tag{2.9}$$

where $\theta$ is the probability that banker $j$ remains a banker in the next period. Inserting the evolution of bankers' equity (2.6) into (2.9) yields:

$$\begin{aligned} V_{jt} = & E_t \sum_{i=0}^{\infty} (1-\theta)\theta^i \beta^{i+1} \Lambda_{t,t+i+1} \\ & \Big[(1 + r^K_{t+1})N_{jt} + (r^K_{t+1} - r^D_t)D_{jt} + (r^K_{t+1} - r^{CB}_t)R^{CB}_{jt}\Big]. \end{aligned} \tag{2.10}$$

With positive premia, bankers have an incentive to blow up their balance sheets infinitely. Following Gertler and Karadi (2011), we introduce a moral hazard to counteract this behavior. Each period, banker $j$ can choose to 'run away', thereby diverting fraction $\lambda$ of the total intermediated funds $Q_t S_{jt}$. In case of such a run, this fraction is lost for households and the central bank.[25] The banker decides to run if income from diverting funds exceeds the expected terminal wealth $V_j$ from being a banker. Hence, $j$'s incentive constraint can be expressed in the following manner:

$$V_{jt} \geq \lambda Q_t S_{jt}. \tag{2.11}$$

Note that banker $j$'s terminal wealth can be expressed recursively as

$$V_{jt} = mu_t^N N_{jt} + mu_t^D D_{jt} + mu_t^R R_{jt}^{CB}. \tag{2.12}$$

The $mu$ variables can be interpreted as the marginal utilities of changes in the different sources of funds:

$$mu_t^N = E_t[(1-\theta)\beta\Lambda_{t,t+1}(1+r_{t+1}^K) + \beta\Lambda_{t,t+1}\theta\Delta_{t,t+1}^N mu_{t+1}^N]; \tag{2.13}$$

$$mu_t^D = E_t[(1-\theta)\beta\Lambda_{t,t+1}(r_{t+1}^K - r_t^D) + \beta\Lambda_{t,t+1}\theta\Delta_{t,t+1}^D mu_{t+1}^D]; \tag{2.14}$$

$$mu_t^R = E_t[(1-\theta)\beta\Lambda_{t,t+1}(r_{t+1}^K - r_t^{CB}) + \beta\Lambda_{t,t+1}\theta\Delta_{t,t+1}^R mu_{t+1}^R], \tag{2.15}$$

where $\Delta_{t,t+1}^N$, $\Delta_{t,t+1}^D$, and $\Delta_{t,t+1}^R$ are the growth rates of equity, deposits, and central bank funds, respectively. Note that we eliminate the $j$ subscripts by assuming that deposits and central bank funds are allocated to banks in accordance with their equity shares — that is $D_{jt} = D_t N_{jt}/N_t$ and $R_{jt}^{CB} = R_t^{CB} N_{jt}/N_t$. Hence, we can derive the growth rates in the following manner:

$$\Delta_{t,t+1}^N = \frac{N_{jt+1}}{N_{jt}} = (1+r_{t+1}^K) + (r_{t+1}^k - r_t^D)\frac{D_t}{N_t} + (r_{t+1}^k - r_t^{CB})\frac{R_t^{CB}}{N_t}; \tag{2.16}$$

$$\Delta_{t,t+1}^D = \frac{D_{jt+1}}{D_{jt}} = \frac{D_{t+1}}{D_t}\Delta_{t,t+1}^N \frac{N_t}{N_{t+1}}; \tag{2.17}$$

$$\Delta_{t,t+1}^R = \frac{R_{jt+1}^{CB}}{R_{jt}^{CB}} = \frac{R_{t+1}^{CB}}{R_t^{CB}}\Delta_{t,t+1}^N \frac{N_t}{N_{t+1}}. \tag{2.18}$$

[25] In reality, banks cannot divert central bank money, as this money is backed by collateral. Thus, for banks, it is not possible to receive additional central bank funds without owning sufficient collateral. Our modeling approach does not imply that bankers will actually ever divert central bank money. Instead, it creates an upper bound for central bank refinancing based on bankers' equity and households' deposits. Thus, we capture banks' natural limits in the acquisition of central bank money, e.g., resulting from insufficient collateral, in a substantially simplified manner.

Inserting (2.12) in (2.11) yields the following incentive constraint:

$$mu_t^N N_{jt} + mu_t^D D_{jt} + mu_t^R R_{jt}^{CB} \geq \lambda Q_t S_{jt}. \tag{2.19}$$

Assuming that the incentive constraint (2.19) is binding and summing across all bankers, we calculate the maximum amount of external refinancing $F^*$:

$$F_t^* = \frac{\lambda - mu_t^N}{mu_t^R - \lambda} N_t + \frac{mu_t^R - mu_t^D}{mu_t^R - \lambda} D_t. \tag{2.20}$$

Accordingly, we express bankers' individual balance sheets (2.5) in aggregate terms in the following manner:

$$Q_t S_t = N_t + D_t + R_t^{CB}. \tag{2.21}$$

Note that $N$ comprises the equity of existing bankers ($N_e$) of new bankers ($N_n$):

$$N_t = N_{et} + N_{nt}. \tag{2.22}$$

$N_e$ can be expressed in the following manner:

$$N_{et} = \theta \Delta_{t-1,t}^N N_{t-1}. \tag{2.23}$$

New bankers receive a fraction $\omega/(1-\theta)$ of the current value of last period's total intermediated funds $Q_t S_{t-1}$. The equity of new bankers can be expressed in the following manner:

$$N_{nt} = \frac{\omega}{1-\theta}(1-\theta) Q_t S_{t-1} = \omega Q_t S_{t-1}. \tag{2.24}$$

### 2.2.3 Intermediate Goods Producers

Intermediate goods producers receive funds exclusively from banks, buy capital goods, and use these capital goods, combined with labor, to produce intermediate goods. Intermediate goods are sold to final goods producers that repackage the intermediate goods and offer them on the goods market. In detail, intermediate goods producers sell $S$ claims to banks at a price $Q$ to obtain funds in return. At the end of period $t$, intermediate goods producers use all the acquired funds to finance investments — that is they buy capital goods $K$ at a price $Q$ per unit. In period

$t+1$, these capital goods are used for production. Consequently, total intermediated funds pose a restriction on the accumulation of capital goods for production.

Following Gertler and Karadi (2011), the price of capital is equal to the price of claims. Therefore, we can express the following equation:

$$Q_t K_{t+1} = Q_t S_t. \tag{2.25}$$

Intermediate goods production is given by the following Cobb-Douglas function:

$$Y_t^M = A_t (U_t \xi_t K_t)^\alpha L_t^{1-\alpha}, \tag{2.26}$$

where $A$ is technology, $U$ the utilization rate of capital, and $\xi$ the quality of capital. Maximizing the profits of intermediate goods producers yields the following first-order conditions for the utilization rate (2.27) and labor demand (2.28):

$$P_t^M \alpha \frac{Y_t^M}{U_t} = \delta'(U_t) \xi_t K_t, \tag{2.27}$$

$$P_t^M (1-\alpha) \frac{Y_t^M}{L_t} = W_t, \tag{2.28}$$

where $P^M$ is the price of intermediate goods and $\delta(U)$ the depreciation rate of capital, with $\delta(U) = \delta_c + U_t^{1+\zeta} b/(1+\zeta)$; $\delta_c$, $b$, and $\zeta$ are adjustment parameters. As all profits from intermediate goods producers are transferred to banks, $R_t^K$ can be written as:

$$R_t^K = \frac{[P_t^M \alpha \frac{Y_t^M}{\xi_t K_t} + Q_t - \delta(U_t)]\xi_t}{Q_{t-1}}. \tag{2.29}$$

Note that the quality of capital ($\xi$) directly affects banks' return on capital. Hence, a negative shock to $\xi$ can induce substantial loan defaults and critical deterioration of banks' balance sheets, which are characteristics of, e.g., the global financial crisis.

### 2.2.4 Capital Goods Producers

Capital goods producers create new capital goods and refurbish depreciated capital goods. The refurbishment cost is fixed at 1, while new capital goods are priced $Q$. The creation of new capital goods is subject to (flow) adjustment costs. Capital producers' profits are transferred in each period to their owners. Gross capital goods created are defined as $I$ and net investment $I^N$ as the difference between $I$ and refurbished capital goods $I^N = I - \delta(U)\xi K$. $\bar{I}$ denotes the steady state level of investment. Capital goods producers maximize the sum of their discounted profits:

$$\max E_t \sum_{i=0}^{\infty} \beta^i \Lambda_{t,t+i} \left[ (Q_{t+i} - 1) I^N_{t+i} - f\Big( \frac{I^N_{t+i} + \bar{I}}{I^N_{t-1+i} + \bar{I}} \Big) (I^N_{t+i} + \bar{I}) \right], \tag{2.30}$$

where $f(\cdot)$ is defined as $\frac{\eta_i}{2} \left[ \frac{I^N_t + \bar{I}}{I^N_{t-1} + \bar{I}} - 1 \right]^2$ with $\eta_i$ as a scaling parameter. Maximizing profits yields the following equation:

$$Q_t = 1 + f(\cdot) + \Big( \frac{I^N_t + \bar{I}}{I^N_{t-1} + \bar{I}} \Big) f'(\cdot) - E_t \beta \Lambda_{t,t+1} \Big( \frac{I^N_{t+1} + \bar{I}}{I^N_t + \bar{I}} \Big)^2 f'(\cdot). \tag{2.31}$$

Hence, in the steady state $\bar{Q} = 1$. Changes in the level of investment increase production costs and, consequently, the price of capital. Note that capital evolves according to the following equation:

$$K_{t+1} = \xi_t K_t + I^N_t. \tag{2.32}$$

### 2.2.5 Final Goods Producers

Final goods producers buy intermediate goods, repackage them, and sell them on the goods market — that is one unit of intermediate goods is converted into one unit of final goods. Final goods producers act as profit-maximizing competitive monopolists. With $\varepsilon$ being the elasticity of substitution, the total output $Y$ is defined as a constant elasticity of substitution (CES) composite of differentiated final goods:

$$Y_t = \left[ \int_0^1 {Y_{ft}}^{\frac{\varepsilon - 1}{\varepsilon}} df \right]^{\frac{\varepsilon}{\varepsilon - 1}}. \tag{2.33}$$

Consumers' cost minimization yields the following definitions for firm $f$'s production $Y_f$ and for prices $P$:

$$Y_{ft} = \left( \frac{P_{ft}}{P_t} \right)^{-\varepsilon} Y_t, \tag{2.34}$$

$$P_t = \left[ \int_0^1 {P_{ft}}^{1-\varepsilon} df \right]^{\frac{1}{1-\varepsilon}}. \tag{2.35}$$

Following Calvo (1983), only the fraction $1-\gamma$ of final goods producers can adjust retail prices in period $t$ to the new optimal level $P^*$. The fraction $\gamma$ of final goods producers is not able to adjust prices to the new optimal level but applies last period's inflation rate $\pi_{t-1,t} = P_t/P_{t-1}$ weighted by an indexation parameter $\gamma_\pi$. Final goods producers do not know, *ex ante*, whether they are able to adjust their prices in the next period. They set prices optimally taking this uncertainty into account. As the only cost factor for final goods producers is the price of intermediate goods $P^M$, their maximization problem can be expressed in the following manner:

$$\max E_t \sum_{i=0}^{\infty} \gamma^i \beta^i \Lambda_{t,t+i} \Big[ \frac{P_t^*}{P_{t+i}} \prod_{k=1}^{i} (\pi_{t+k-1,t+k})^{\gamma_\pi} - P_{t+1}^M \Big] Y_{ft+i}. \tag{2.36}$$

Applying the law of large numbers yields the following definition of retail prices:

$$P_t = [(1-\gamma)(P_t^*)^{1-\varepsilon} + \gamma(\pi_{t-1,t}^{\gamma_\pi} P_{t-1})^{1-\varepsilon}]^{\frac{1}{1-\varepsilon}}. \tag{2.37}$$

Thus, the retail price level is a weighted average of adjusted and non-adjusted prices.

### 2.2.6 Central Bank

The central bank sets the nominal interest rate on central bank funding $i^{CB}$ according to a standard Taylor rule without interest rate smoothing (Gertler and Karadi, 2011). Interest rates on different forms of saving — bonds, CBDC, and bank deposits — depend on $i^{CB}$ to ensure that $i^B \geq i^D \geq i^{CBDC}$ (see Table 2.1). In this manner, the central bank 'leads' all interest rates with its rule-based interest rate on central bank funding:

$$i_t^{CB} = (1 + \bar{r}^{CB}) + \kappa_\pi \pi_t + \kappa_{y_{gap}} y_{gap,t}, \tag{2.38}$$

where $\kappa_\pi$ is the inflation weight, $\kappa_{y_{gap}}$ the weight of the output gap, and $\bar{r}^{CB}$ the neutral (steady state) real interest rate. Following Gertler and Karadi (2011), we use minus the price markup as a proxy for the output gap.

We assume that the nominal interest rate on deposits follows the interest rate on central bank funding with the fixed spread $\Delta^D$:[26]

$$i_t^D = i_t^{CB} - \Delta^D. \tag{2.39}$$

We introduce this spread to match data indicating that, in normal times, central bank refinancing is more expensive than refinancing via deposits (Bindseil, 2020). While a fixed spread is a simplified assumption, it is heavily used in the literature, e.g., in Bindseil (2020).

In Chapter 2.4, we analyze scenarios, in which the ELB is binding. In these cases, if the interest rate on deposits would become negative, it is constrained by the ELB.[27] Accounting for the ELB, the interest rate on deposits is determined as follows:

$$i_t^D = \begin{cases} i_t^{CB} - \Delta^D & \text{for } i_t^{CB} - \Delta^D \geq 0, \\ 0 & \text{for } i_t^{CB} - \Delta^D < 0. \end{cases} \tag{2.40}$$

The central bank also sets the interest rate on CBDC. We explicitly differentiate between an interest-bearing CBDC and a non-interest-bearing CBDC. In the case of a non-interest-bearing CBDC, we set $i^{CBDC}$ to zero:

$$i_t^{CBDC} = 0. \tag{2.41}$$

For an interest-bearing CBDC, the interest rate on CBDC strictly follows the interest rate on central bank funding with the fixed spread $\Delta^{CBDC}$, such that $i^{CBDC} < i^{CB}$, as proposed in Bindseil (2020):

$$i_t^{CBDC} = i_t^{CB} - \Delta^{CBDC}. \tag{2.42}$$

In Chapter 2.4.4, we decouple these interest rates and allow for an individual rule-based determination, in which the CBDC rate is used as a policy tool. Note that the interest rate on CBDC can be negative.

The interest rate on government bonds follows the interest rate on central bank funding with the fixed spread $\Delta^B$. We assume a positive spread based on bond yield data for the period before the global financial crisis and the rationale that the

[26] Note that in reality, banks determine the interest rate on deposits themselves. However, maximizing their profits, banks use the central bank-set interest rates as the benchmark rate, as indicated by a high correlation between these interest rates.

[27] In the following, we assume that the ELB lies at 0% and is, thus, a zero lower bound.

lack of liquidity services has to be compensated for by a higher remuneration.[28]

$$i_t^B = i_t^{CB} + \Delta^B. \tag{2.43}$$

The connection between nominal and real interest rates is given by the following Fisher relations:

$$1 + i_t^D = (1 + r_t^D)(1 + E_t\pi_{t,t+1}); \tag{2.44}$$

$$1 + i_t^{CBDC} = (1 + r_t^{CBDC})(1 + E_t\pi_{t,t+1}); \tag{2.45}$$

$$1 + i_t^B = (1 + r_t^B)(1 + E_t\pi_{t,t+1}). \tag{2.46}$$

Apart from setting interest rates, the central bank also provides funding to commercial banks via central bank loans. As refinancing via the central bank is more expensive than refinancing via deposits ($r^{CB} > r^D$), banks will only demand central bank funding ($R^{CB}$) to fill the gap between the supply of deposits ($D$) and the maximum amount of total external refinancing ($F^*$):

$$R_t^{CB} = F_t^* - D_t. \tag{2.47}$$

Note that this expression implicitly assumes a full allotment procedure: As long as the banks' incentive constraint holds — that is, as long as they can provide sufficient collateral —, the central bank fully meets their money demand. We relax this assumption of full allotment in Chapter 2.4.3.

### 2.2.7 Government and Aggregation

The government receives income from lump-sum taxes $T$ and issues government bonds $B_t$. It finances government spending ($G$) and repays last period's bond holdings $B_{t-1}$ including interest payments $i_{t-1}^B$. Note that we define $G$ as a constant share of steady state output.

$$\bar{G} + (1 + i_{t-1}^B)B_{t-1} = T + B_t. \tag{2.48}$$

[28]Note that the fixed spread is a simplifying assumption. In reality, bond prices and yields exhibit more complex dynamics.

Output is divided into consumption, investment, investment adjustment costs, and government expenditures. Hence, the economy-wide budget constraint can be expressed in the following manner:

$$Y_t = C_t + I_t + f\Big(\frac{I_t^N + \bar{I}}{I_{t-1}^N + \bar{I}}\Big)(I_t^N + \bar{I}) + \bar{G}. \tag{2.49}$$

## 2.3 Calibration

Table 2.2 summarizes the calibration of our model. We use a total of 24 parameters, 17 of which are conventional and also used in Gertler and Karadi (2011). We introduce additional parameters related to the inclusion of money in the utility function ($\Upsilon$, $\Gamma$), the discount factor $\psi$ ($\Omega_D$, $\Omega_N$), and the interest rate spreads ($\Delta^B$, $\Delta^D$, $\Delta^{CBDC}$). Since no CBDC has been introduced in an industrialized economy thus far, there is a lack of micro data for the key parameters related to CBDC. Therefore, we calibrate these parameters to match available macro data in the absence of CBDC.

The calibration of the conventional parameters closely follows that of Gertler and Karadi (2011). Our calibration differs in terms of the following two aspects: First, we derive the discount factor $\beta$ from the data for the average bond interest rate from 2003 to 2008 (Bindseil, 2020)). Second, we adjust the steady state government expenditure share to match euro area data (Eurostat, 2020).

We calibrate the additional parameters in the following manner. We use $\Omega_D$ to target a steady state share of central bank funding of 17% in external refinancing.[29] Note that, due to the functional form of $\psi$, higher values for $\Omega_D$ do not only decrease the aforementioned share but also the elasticity of households' deposits to changes in interest rates. $\Omega_N$ is used to define the impact of financial stress on deposits. As there is no reliable euro area data on how households adjust their bank deposits in times of financial crisis and in the absence of deposit insurance schemes, we calibrate $\Omega_N$ such that — with CBDC — deposits initially drop approximately by 20% after the shock. $\Upsilon$ and $\Gamma$ determine the absolute and the marginal utility of liquidity,

[29] From 2003–2008, central bank refinancing, on average, accounted for 3% of bank funding, while capital market refinancing accounted for 30% (Bindseil, 2020). In our analysis, we neglect capital market refinancing. As a consequence, it seems reasonable to assume a higher share of central bank funding than the 3% outlined in Bindseil (2020).

**Table 2.2:** Parameter calibration

| | | |
|---|---|---|
| **Households** | | |
| $\beta$ | Intertemporal Discount Factor | 0.990 |
| $h$ | Habit Parameter for Consumption | 0.815 |
| $\chi$ | Relative Utility Weight of Labor | 3.409 |
| $\phi$ | Inverse Frisch Elasticity of Labor Supply | 0.276 |
| $\Upsilon$ | Utility Weight of Liquidity | 0.125 |
| $\Omega_D$ | Elasticity of $\psi$ to Bank Deposits | 51.000 |
| $\Omega_N$ | Impact of Financial Stress on $\psi$ | 0.050 |
| $\Gamma$ | Elasticity of Liquidity | $-0.950$ |
| **Banks** | | |
| $\theta$ | Survival Probability of Bankers | 0.975 |
| $\lambda$ | Divertible Fraction of Intermediated Funds | 0.381 |
| $\omega$ | Proportional Transfer to Entering Bankers | 0.002 |
| **Intermediate Goods Producers** | | |
| $\alpha$ | Capital Share | 0.330 |
| $\zeta$ | Elasticity of Marginal Depreciation | 7.200 |
| $\delta_i$ | Steady State Depreciation Rate | 0.025 |
| **Capital Goods Producers** | | |
| $\eta_i$ | Elasticity of Investment Adjustment Costs | 1.728 |
| **Final Goods Producers** | | |
| $\varepsilon$ | Elasticity of Substitution between Goods | 4.167 |
| $\gamma$ | Calvo Parameter | 0.779 |
| $\gamma_\pi$ | Price Indexation of Inflation | 0.241 |
| **Central Bank and Government** | | |
| $\kappa_\pi$ | Taylor Rule Response Coefficient to Inflation | 1.500 |
| $\kappa_{y_{gap}}$ | Taylor Rule Response Coefficient to Output Gap | 0.5/4 |
| $\Delta^B$ | Spread between Central Bank Reserves and Bonds | 0.01/4 |
| $\Delta^D$ | Spread between Central Bank Reserves and Deposits | 0.01/4 |
| $\Delta^{CBDC}$ | Spread between Central Bank Reserves and CBDC | 0.02/4 |
| $\bar{G}/\bar{Y}$ | Steady State Share of Government Expenditures | 0.470 |

respectively. We calibrate both parameters such that households do not hold any non-interest-bearing CBDC in the steady state — that is households' bank deposits fully meet their liquidity needs.

The model features four different interest rates. In the baseline setting, we assume that $r^D$, $r^B$, and $r^{CBDC}$ follow $r^{CB}$ with time-invariant spreads. $\Delta^B$ and $\Delta^D$ are set to 1%, such that $\bar{r}^B = 4\%$ and $\bar{r}^D = 2\%$ approximately match the observed data. Following Bindseil (2020), we assume that in the steady state, the CBDC rate lies 2% below the interest rate on central bank loans. As the model output presents quarterly data, interest rate spreads are adjusted accordingly.

## 2.4 Introducing CBDC

In this chapter, we discuss the implications of two different forms of CBDCs, an interest-bearing and a non-interest-bearing CBDC. For an interest-bearing CBDC, the central bank sets a flexible interest rate that can be either positive or negative. In contrast, a non-interest-bearing CBDC is not remunerated and is, in this respect, the digital equivalent of cash. In a cashless economy, these two CBDC alternatives differ fundamentally: a non-interest-bearing CBDC anchors interest rates and imposes, just like cash, an ELB on deposit interest rates. The interest-bearing alternative imposes a similar lower bound. However, this lower bound can be flexible and co-moves with the CBDC interest rate.[30] Therefore, the central bank can react to a crisis by setting interest rates below the original ELB — that is, in our case, below zero — and stimulate the economy more effectively.

Our CBDC analysis involves four steps: First, in Chapter 2.4.1, we compare the baseline model without CBDC with a non-interest-bearing CBDC model under the impact of a quality of capital shock. We assume that both models are constrained by an ELB. Second, in Chapter 2.4.2, we use the same shock to compare the baseline model (ELB-constrained and -unconstrained) to an unconstrained interest-bearing CBDC model. Third, in Chapter 2.4.3, we relax the assumption of full allotment of central bank money. Finally, in Chapter 2.4.4, we conclude with an analysis of a flexible rule-based interest rate on CBDC, such that the CBDC interest rate is used

[30]Note that this variability of the lower bound only holds in a cashless society, which we assume for our analysis.

as an additional monetary policy tool.

We choose this order, as it allows us to address CBDC implications step-by-step. The first two sections highlight the reallocation of households' savings and the resulting change in the structure of bank funding. These sections also establish the general result that full allotment can replace losses in bank funding and offset negative consequences beyond the financial sector. Relaxing the assumption of full allotment, we first focus on the impact of a CBDC on the real economy and, then, on the central bank's option to use the interest rate on CBDC as an additional monetary policy tool to mitigate destabilizing effects.

For all simulations, we use a negative quality of capital shock of 5% with persistence 0.66 to simulate a financial crisis that features substantial loan defaults, such that the simulation leads to dynamics comparable to the global financial crisis (Gertler and Karadi, 2011). The general model mechanics and a comparison to Gertler & Karadi's model is presented in Chapter 2.6.2.[31]

### 2.4.1 Non-Interest-Bearing CBDC

Figure 2.3 compares the dynamics of the baseline model without a CBDC with a model with a non-interest-bearing CBDC. The negative quality of capital shock implies a major reduction in the output of intermediate goods. This reduction leads to loan defaults[32] and a deterioration of banks' balance sheets.

A 5% quality of capital shock amounts to a default of approximately 70% of loans, thereby resulting in an equally high percentage loss of bank equity. The starting recession and deflationary developments call the central bank into action. The central bank lowers the nominal interest rate on central bank funding to stimulate lending and investment. Accordingly, also the interest rate on deposits drops. As the non-interest-bearing CBDC imposes an ELB, the deposit interest rate remains slightly above the CBDC interest rate.

[31]We conduct our simulations using Dynare (Adjemian et al., 2011) and implement occasionally binding constraints via OccBin (Guerrieri and Iacoviello, 2015). We provide additional impulse response functions (IRFs) for additional variables in Chapter 2.6.3.

[32]Note that there are no actual loan defaults in the model. The fall in capital efficiency leads to a fall in firm value and, hence, in bank equity because banks are the residual owners of firms. Following Gertler and Karadi (2011), this mechanism can be broadly interpreted as a loan write-off.

**Figure 2.3:** Baseline with ELB vs. non-interest-bearing CBDC with ELB

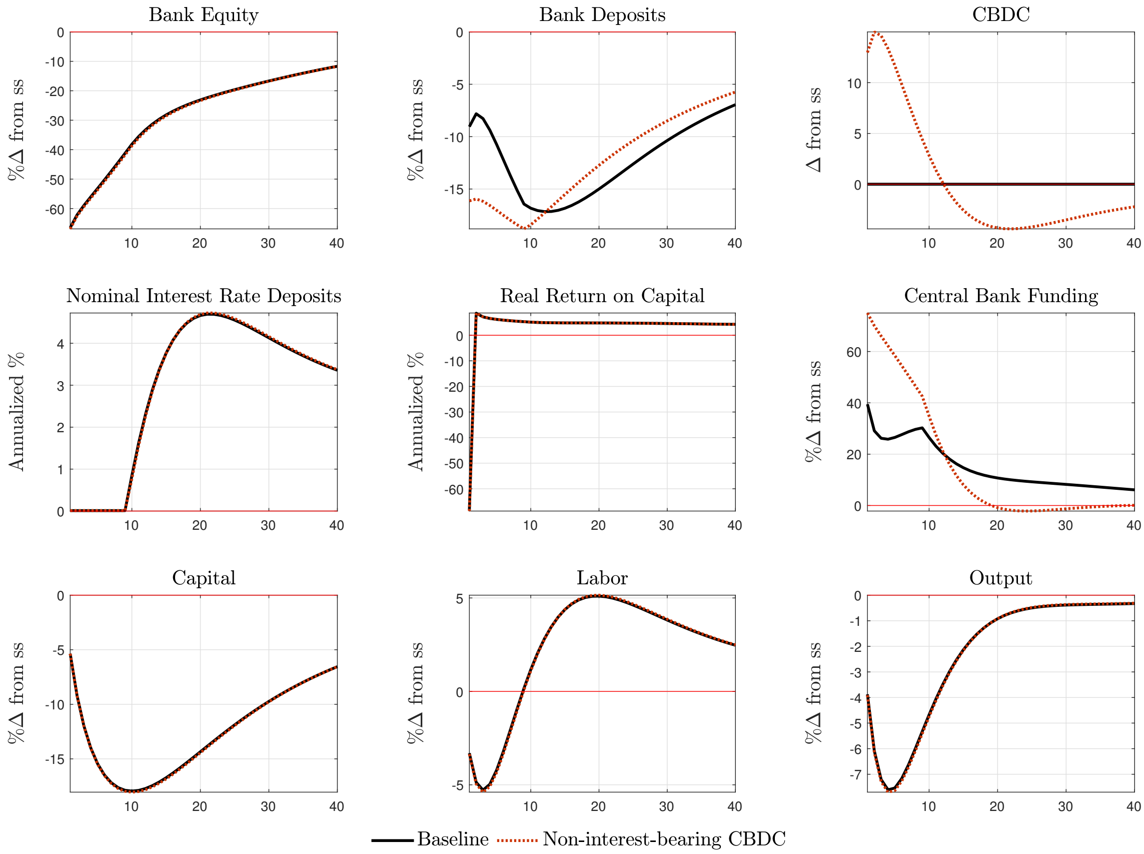


The lower spread between bank deposits and CBDC incentivizes households to substitute bank deposits with CBDC. Based on our calibration, with CBDC, bank deposits decrease by an additional 7%. This reduction in deposits leads to an increase in central bank funding by 70%, as banks substitute lost funds from households with central bank funds. The share of central bank funds in the external refinancing of banks increases from initially 17% to 29%.

The central bank's balance sheet is additionally extended, in the case with a CBDC, as households deposit their savings with the central bank — that is in CBDC. Note that the main reason for the substantial increase in CBDC is not the decline in deposits. Instead, as the interest rate on bonds declines, households, additionally, substitute bonds with CBDC. This effect is in line with the observed increased use of central bank money (cash) in times of financial distress. As a CBDC offers the same attractive features as cash — a constant, non-negative, and guaranteed nominal interest rate of zero — but imposes no marginal costs, a non-interest-bearing CBDC might be used intensively as a store of value in times of low interest rates.[33] As the economy recovers and prices rise above the steady state level, the central bank reacts by increasing the interest rate on central bank funding. Accordingly, the deposit interest rate follows, and the spread between CBDC and alternative forms of savings increases. As the effect overshoots steady state levels, households decrease their CBDC holdings below zero.[34] Part of the liquidity created by CBDC debt is deposited with banks, where households profit from the increased spread, such that bank deposits in the CBDC model exceed their counterpart in the baseline model after period twelve. With the increase in bank deposits, central bank funds slowly return to the steady state level.

There are only minor effects on refinancing and production. First, banks rely more on central bank funding. Therefore, they initially face lower refinancing costs, as the interest rate on central bank funding is not constrained by an ELB. As interest rates quickly recover in the first 10 periods and central bank funds are reduced, this effect is relatively small. Second, as households substitute CBDC for bank deposits, they experience a change in their budget constraint, thereby leading to a small reduction in labor supply — and thus output — of further 0.05%.

To summarize, the major effects of a non-interest-bearing CBDC are limited to the financial sector and do not substantially affect production. Any losses in deposits are counterbalanced by a one-to-one increase in central bank funds. Thus, losses in deposits do not affect total intermediated funds, as the size of bank's balance sheets does not change. Hence, capital does not deviate from its baseline path, thereby creating no further disturbances in labor, output, and real return on intermediated funds. Note that this neutrality is driven by the assumption of full allotment. This result is in line with Brunnermeier and Niepelt (2019) and Niepelt (2020).

[33] In this simulation, CBDC deposits increase substantially and exceed central bank funds provided to banks by a factor of 6.5, thereby leading to a considerable expansion of the central bank's balance sheet. Considering that, according to Eurostat and ECB data, the total net financial assets of households in the euro area amount to approximately 34,000 billion euro and central bank reserves that account for 3% of banks' external refinancing amount to approximately 624 billion euro, this value seems high but not implausible.

[34] Note that the negative values of CBDC can occur due to technical limitations of the OccBin toolbox. However, in the subsequent analyses, we impose an occasionally binding constraint and prevent negative values of CBDC.

### 2.4.2 Interest-Bearing CBDC

Figure 2.4 depicts the simulation results for the baseline model with and without an ELB and a model with an interest-bearing CBDC.[35] We present the baseline model both with and without an ELB to highlight that the major real effects do not occur due to disturbances caused by the CBDC. Instead, the real effects can be explained by the circumvention of the ELB. We assume that, in the CBDC model, households do not have access to cash or any other non-interest-bearing asset. Hence, there is no way to avoid negative interest rates, and the ELB is no longer imposed, thereby allowing deposit interest rates to below zero.

**Figure 2.4:** Baseline with ELB vs. interest-bearing CBDC

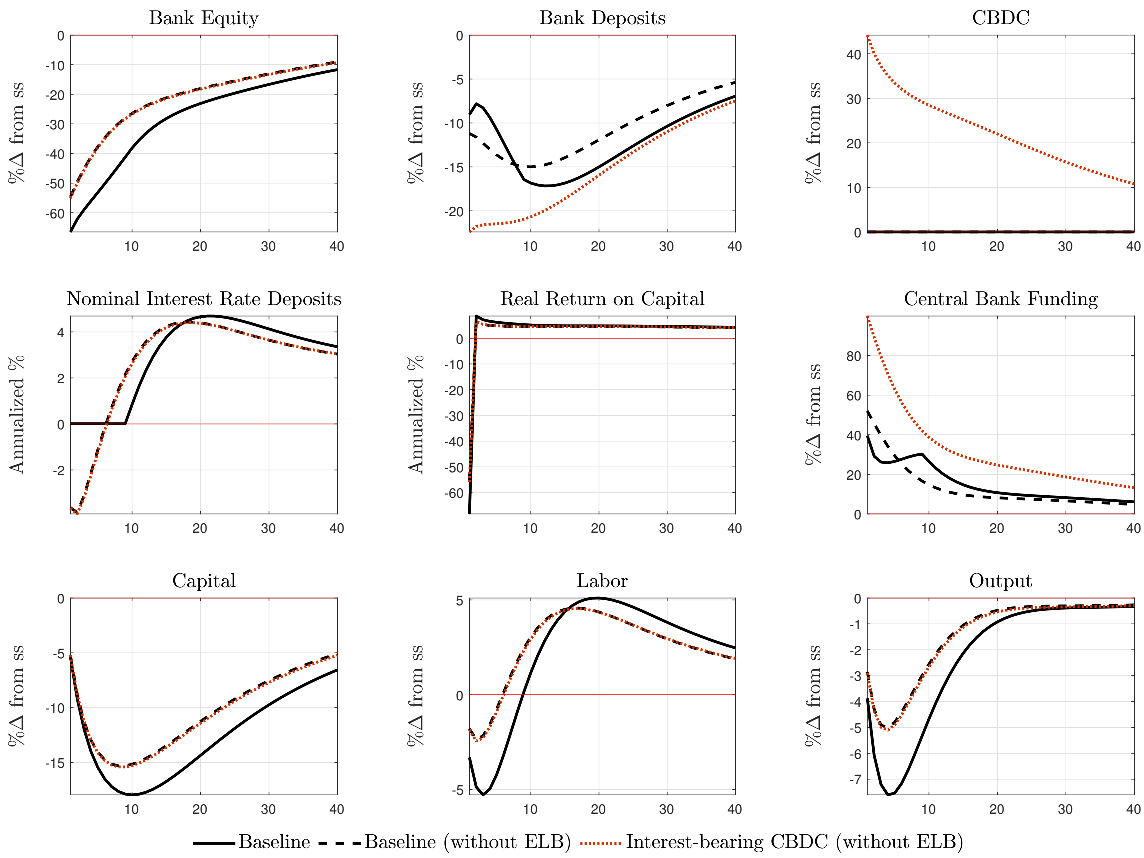


The major advantage of an unconstrained deposit interest rate for monetary policy is that monetary policy measures directly affect households' savings decisions, for

[35]We acknowledge that negative interest rates on CBDC are controversial. In this paper, we do not address associated concerns, but solely focus on monetary policy aspects.

positive as well as negative interest rates. In this case, the nominal deposit interest rate follows the interest rate on central bank funds set by the central bank based on the Taylor rule. Hence, the central bank's reaction to economic changes — that is the inflation rate and the output gap — translates directly to households. Lower deposit interest rates incentivize households to initially increase labor by approximately 1.5% and lead to a 1% higher output compared to the ELB-constrained baseline model. In addition, lower deposit interest rates imply a higher premium for banks and accelerate the build-up of new equity. Therefore, in the unconstrained case, monetary policy is better equipped to mitigate adverse effects. The stronger reduction in the nominal interest rate on bank deposits leads to a further decline in deposits by 2%. This decline becomes larger and moves to 11% when households have the opportunity to shift savings to an equally liquid CBDC. Note that this effect is not driven by changes in the interest rate spread. Instead, as financial stress reduces households' demand for deposits, a CBDC offers a viable alternative to satisfy their demand for liquidity. By holding CBDC, households increase their overall liquidity, while the marginal utility of liquidity decreases. This *liquidity effect* renders deposits less attractive and leads to a further reduction in deposits.[36] In the steady state, households hold approximately 27% of their liquidity in CBDC.[37] Initially, after the shock, this share increases to 41%. Simultaneously, the loss in deposits is offset by an increase in central bank funds. The share of central bank funding in total external refinancing doubles from 18% to 36%. In contrast to the non-interest-bearing CBDC model, CBDC only slightly exceeds central bank funds in the central bank's balance sheet ($CBDC/R^{CB} = 1.25$).

Again, for the same reasons discussed in the previous section, the major effects of the interest-bearing CBDC are limited to the financial sector and do not substantially affect production. However, taking into account that an interest-bearing CBDC might eliminate the ELB, it improves the monetary policy transmission and enables the central bank to counteract a financial crisis more efficiently. Nevertheless, this effect on the real economy, including production, is not directly linked to CBDC or

[36] Note that this drop is additionally amplified by a comparably high elasticity of demand for deposits on changes in banks' equity.

[37] This value results from two assumptions. First, in the steady state, the remuneration for CBDC is 1%. Second, for consistency, we apply the same parametrization (particularly $\Upsilon$) as in the non-interest-bearing CBDC model.

changes in the households' saving options, but the elimination of the ELB. Note that, again, these results are driven by the assumption of full allotment. This assumption is relaxed in the next section.

### 2.4.3 Alternative Allotment of Central Bank Funds

Thus far, we assumed that the central bank fully compensates for losses in deposits by providing additional central bank funds. This assumption is in line with the current monetary policy of the ECB that, as a reaction to the global financial crisis, adapted its tender procedure for open market operations to full allotment in October 2008. The ECB began to fully allocate demanded funds to banks to stabilize the interbank market. While full allotment currently appears to be the 'new normal', it should not be taken for granted.

This observation begs the question of whether our results still hold under alternative allotment procedures. In fact, as we show in this section, the assumption of full allotment is necessary to obtain the result that CBDC does not affect the economy beyond the financial sector.

To analyze restricted allotment, we adapt 2.47 in the following manner:

$$R_t^{CB} = \bar{R}^{CB} + X[(F_t^* - \bar{F}^*) - (D_t - \bar{D})], \tag{2.50}$$

where $X$ is the share of lost deposits outside the steady state that the central bank substitutes. Thus, losses of deposits after a shock are only partially compensated. Note that this functional form does not affect the steady state allocation of central bank funds, such that $\bar{R}^{CB}$ is equal in all models. Thus, the results from different model specifications are comparable.

Figure 2.5 compares the baseline model for full allotment and restricted allotment ($X = 0.5$) with the interest-bearing CBDC model ($X = 0.5$). All models are not constrained by the ELB. Note that the central bank decides on the fraction of compensated funds. The more funds the central bank provides, the lower the real effects. In our simulation, we use $X = 0.5$ as an example.

**Figure 2.5:** Interest-bearing CBDC with different allotment of central bank funds

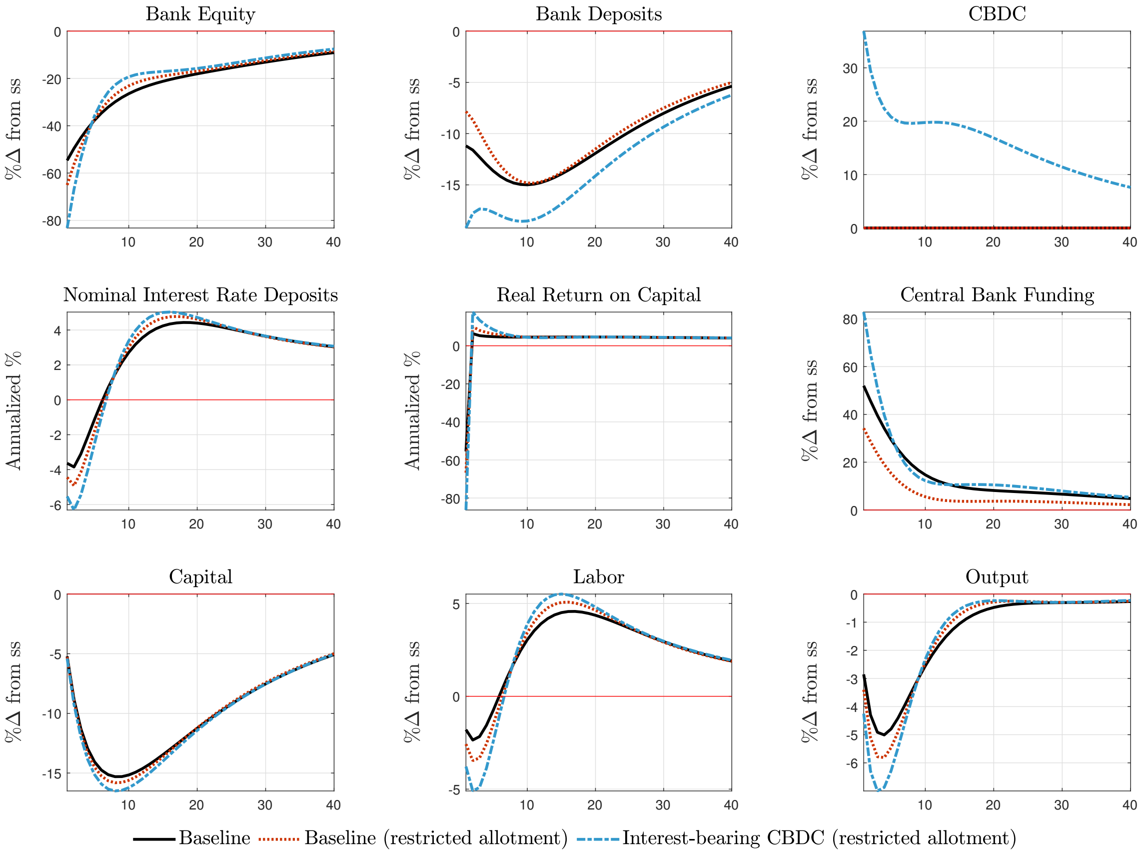


As the central bank does not fully compensate for lost deposits in both models, total intermediated funds, and, thus, the size of banks' balance sheets, decrease. This decrease negatively affects the next periods' levels of capital, thereby resulting in lower output. In addition, lower levels of capital increase the marginal productivity of capital and decrease the marginal productivity of labor. Hence, the real return on capital increases in periods after the initial shock while wages drop. Households react with a reduction in labor, which is, due to consumption smoothing, already present in the first period. With $X = 0.5$, this 0.5% stronger drop in labor results in a 0.3% lower output in the baseline model. In the interesting-bearing CBDC model, labor drops an additionally 2%, leading to a further decline in output by 1.2%. The real return on capital and, thus, banks' equity drop an additional 10% in the baseline model and 25% in the interest-bearing CBDC model. The central bank reacts with a reduction in interest rates. This reduction, in combination with the higher expected return on capital, increases the premium and profits for banks. As these higher expected profits ease the moral hazard problem, households are willing to deposit more funds with banks. Even though this easing increases the central bank's willingness to provide funds, central bank funding decreases due to the lower allotment rate. Driven by the high premia, banks promptly restore large parts of their equity and trigger an accelerated recovery process for the entire economy.

With CBDC, households have an incentive to exchange parts of their deposits for CBDC. Thus, deposits and total intermediated funds as well as capital decrease. As described above, this decrease further eases the moral hazard problem, and the central bank provides more funds. Nevertheless, this increase in central bank funding cannot fully compensate for the increased loss in deposits, thereby leading to a deeper recession.

In summary, generalizing the assumption of full allotment leads to remarkably different results. The resulting imperfect substitution of deposits with central bank funds opens up a channel for CBDC to the real economy. The disintermediation of commercial banks negatively impacts investment, the build-up of capital, and production. In this case, CBDC indeed has the potential to destabilize the financial sector and the entire economy.

### 2.4.4 CBDC Interest Rate Rule

While the previous analysis suggests that full allotment is necessary to prevent destabilizing effects, the central bank can also use another tool. Bindseil (2020) proposes that central banks can actively use the interest rate on CBDC to disincentivize its accumulation in a crisis and, thus, to counteract disintermediation. Using this new policy instrument, the central bank can try to govern the demand for CBDC. As the CBDC interest rate in our model is close to zero in the steady state, this approach implies negative interest rates.

For the following analysis, we adapt the CBDC interest rate rule (2.42) in the following manner:

$$i_t^{CBDC} = i_t^{CB} - \left(\Delta^{CBDC} + \frac{\bar{N} - N_t}{\bar{N}}\kappa_N\right). \tag{2.51}$$

The term in parentheses defines the spread between the interest rates on central bank funding and CBDC. We keep its steady state level unchanged and allow the central bank to increase the spread based on financial stress after the shock. We use the measure from Chapter 2.2.1, such that financial stress is expressed as the percentage deviation of banks' equity from steady state. $\kappa_N$ specifies the intensity of the reaction.[38]

**Figure 2.6:** Flexible interest rate spread on CBDC

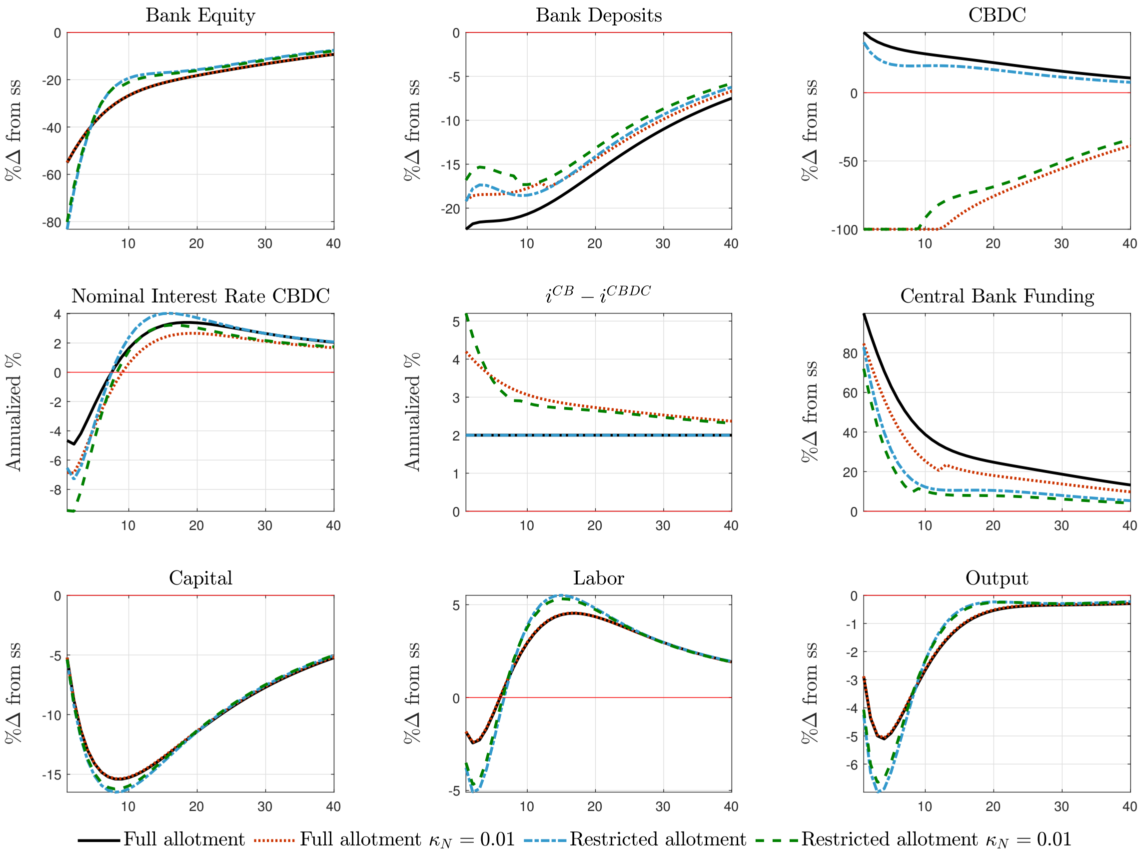


The blue and the green lines in Figure 2.6 indicate the results for models with restricted allotment ($X = 0.5$). As expected, decreasing the nominal interest rate on CBDC reduces CBDC holdings — in our case to zero.

The effect on deposits is relatively small, as households do not substitute CBDC primarily with deposits but with bonds. The liquidity effect drives the smaller drop in deposits: As households decrease their CBDC holdings, total liquidity declines, and its marginal utility rises. This effect increases the marginal utility of deposits,

[38] $\kappa_N$ is calibrated such that households in this exercise initially reduce their CBDC holdings to zero. Note that we restrict these holdings to be non-negative.

and thus, deposits themselves, but is outweighed by the rising risk.[39] With restricted allotment, (relatively) higher deposits increase total intermediated funds and result in higher labor, capital, and output. However, all these improvements fall short of the full allotment scenario. In other words, while targeting CBDC can positively impact an economy with restricted allotment in a crisis, full allotment is the more effective policy. Nevertheless, lowering interest rates effectively limits the accumulation of CBDC and is a valid tool to mitigate disintermediation and destabilization specifically caused by a CBDC.

With full allotment, the CBDC interest rate proves to be an effective instrument to impact both CBDC holdings and central bank funds. When the interest rate is reduced, households decide to hold less CBDC and more deposits, such that the share of central bank funding in total external refinancing decreases. Thus, there is a twofold contraction in the central bank's balance sheet while economic activity is unaffected.

## 2.5 Conclusion

While CBDCs can offer several benefits to individuals, their implications for the financial sector in general and commercial banks' funding, in particular, remain subject to debate. To contribute to this debate, we developed a medium-sized DSGE model that provides a basis for analyzing the effects of CBDCs. The model features endogenously limited bank funding via households and the central bank, households that actively choose the amount of deposits as part of their utility maximization, and a CBDC as a liquidity-providing substitute for deposits. In addition, our model includes specific interest rates on bonds, deposits, central bank funds, and CBDC, and can account for an ELB on nominal interest rates.

The design of the model implies that households reduce their deposits with commercial banks in times of crises due to a liquidity effect. When households can satisfy their demand for liquidity with CBDC, their main incentive to store their savings in the form of risky deposits is mitigated. The resulting disintermediation implies

[39]Note that CBDC is increasingly attractive when deposits fall, such that households almost fully substitute lost liquidity. Vice versa, this is not the case. The attractiveness of deposits only partially depends on the presence or absence of CBDC (liquidity effect). The determining factor is households' perceived risk of commercial bank money. Households are willing to forgo liquidity when remuneration on CBDC is too low to avoid this risk.

a contraction in the balance sheets of commercial banks and, thus, reduced loan volume, investment, and economic activity.

In our model, the central bank has two options to react to this disruption in commercial bank funding and combat destabilizing effects. First, it can adjust its allotment policy. When faced with a decreasing supply of deposits, commercial banks increase their demand for central bank funds. In case the central bank chooses to fully meet this demand, a reduction in deposits only implies a shift in the composition of bank funding but no contraction of banks' balance sheets. The central bank commits to substitute lost deposits with additional central bank funds, thereby substantially expanding its own balance sheet. While we abstract from the aspect of collateral in our model, the question remains whether banks can provide sufficient eligible assets. If collateral is scarce, the central bank might be pressurized to reduce collateral requirements — that is, it might accept collateral with higher risk, potentially threatening financial stability. Further research is needed to address these issues.

Second, the central bank can decrease the remuneration of CBDC to disincentivize its accumulation. This approach effectively lowers CBDC holdings but does not necessarily incentivize households to hold substantially more deposits. Therefore, on its own, it might not be a sufficient tool to counteract the adverse effects resulting from losses in bank funding in a crisis. Nevertheless, lowering interest rates effectively limits the accumulation of CBDC and is a useful tool to mitigate disintermediation and destabilization caused specifically by a CBDC. It helps control the demand of CBDC and central bank funds without causing CBDC-specific disturbances beyond the financial sector. Note that this second option is only available for an interest-bearing CBDC. For a non-interest-bearing CBDC, the central bank cannot directly steer the demand and prevent substantial accumulation. Apart from a strong commitment to full allotment, at least two alternative policies can mitigate CBDC-induced disintermediation. First, the central bank can limit the supply of CBDC, for example, by imposing a cap on individual CBDC holdings, as proposed by Panetta (2018). However, a cap could weaken a CBDC's competitiveness relative to private digital means of payment, such as global stablecoins, reducing one of the key motives for introducing a CBDC. Second, policy-makers could target the perceived risk in the financial sector by providing deposit insurance schemes, such as those implemented in Germany. While these schemes helped to maintain trust in the financial sector during the global financial crisis, there is evidence that deposit insurances themselves can threaten financial stability (Demirgüç-Kunt and Detragiache, 2002). Further research is needed to analyze CBDC in a model that includes deposit insurance schemes.

Apart from the limitations of our analysis mentioned above, two additional aspects are worth pointing out: First, we model government bonds in a rather simplistic manner. We neglect that the supply of bonds could be limited and that prices and yields are determined by supply and demand in capital markets. Increasing collateral needs from commercial banks would affect demand for bonds and might open up new channels for a CBDC to impact the economy even with full allotment. Second, we analyze the impact of a CBDC in a cashless economy. Since, currently, households continue to hold substantial amounts of their savings in cash, a model including cash could provide further relevant insights.

## 2.6 Appendix

### 2.6.1 Households' Maximization Problem

Households maximize their utility based on the following five variables: consumption $C$, labor $L$, bank deposits $D$, central bank digital currency $CBDC$, and government bonds $B$. Households' utility function comprises a standard log-utility from consumption with habit formation, disutility from labor, and utility from liquidity:

$$\max E_t \sum_{i=0}^{\infty} \beta^i \Big( ln(C_{t+i} - hC_{t+i-1}) + \frac{\Upsilon}{1+\Gamma} \ (D_{t+i} + CBDC_{t+i})^{1+\Gamma} - \frac{\chi}{1+\phi} L_{t+i}^{1+\phi} \Big). \tag{2.52}$$

Households' budget constraint can be written in the following manner:

$$C_t + D_t + CBDC_t + B_t = w_t L_t + \Pi_t + (1 + r_{t-1}^D)\psi_{t-1} D_{t-1} + (1 + r_{t-1}^{CBDC}) CBDC_{t-1} + (1 + r_{t-1}^B) B_{t-1}, \tag{2.53}$$

with

$$\psi_t = 1 - \left( \frac{D_t}{F_t^*} \right)^{\Omega_D} - \frac{\bar{N} - N_t}{\bar{N}} \Omega_N. \tag{2.54}$$

To derive households' savings decision, we set up the Lagrangian in the following manner:

$$\begin{aligned}
\mathcal{L} = & E_t \sum_{i=0}^{\infty} \beta^i \Big\{ ln(C_{t+i} - hC_{t+i-1}) + \frac{\Upsilon}{1+\Gamma} \ (D_{t+i} + CBDC_{t+i})^{1+\Gamma} \\
& - \frac{\chi}{1+\phi} L_{t+i}^{1+\phi} \\
& - \lambda_{t+i} \left[ C_{t+i} + D_{t+i} + CBDC_{t+i} + B_{t+i} - w_{t+i} L_{t+i} - \Pi_{t+i} \right. \\
& - (1 + r_{t+i-1}^D)(1 - \left( \frac{D_{t+i-1}}{F_{t+i+1}^*} \right)^{\Omega_D} - \frac{\bar{N} - N_{t+i-1}}{\bar{N}} \Omega_N) D_{t+i-1} \\
& \left. - \ (1 + r_{t+i-1}^{CBDC}) CBDC_{t+i-1} - (1 + r_{t+i-1}^B) B_{t+i-1} \right] \Big\}.
\end{aligned} \tag{2.55}$$

Now, we derive the Lagrangian with respect to $C_t$, $L_t$, $D_t$, $CBDC_t$, and $B_t$:

$$\frac{\partial \mathcal{L}}{\partial C_t} = (C_t - hC_{t-1})^{-1} - \beta h(C_{t+1} - hC_t)^{-1} - \lambda_t; \tag{2.56}$$

$$\frac{\partial \mathcal{L}}{\partial L_t} = -\chi L_t^{\phi} + \lambda_t w_t; \tag{2.57}$$

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial D_t} = {} & \Upsilon(D_t + CBDC_t)^{\Gamma} - \lambda_t \\ & + \beta\lambda_{t+1}(1 + r_t^D)\left\{\psi_t - \Omega_D\left(\frac{D_t}{F_t^*}\right)^{\Omega_D}\right\};\end{aligned} \tag{2.58}$$

$$\frac{\partial \mathcal{L}}{\partial CBDC_t} = \Upsilon(D_t + CBDC_t)^{\Gamma} - \lambda_t + \beta\lambda_{t+1}(1 + r_t^{CBDC}); \tag{2.59}$$

$$\frac{\partial \mathcal{L}}{\partial B_t} = -\lambda_t + \beta\lambda_{t+1}(1 + r_t^B). \tag{2.60}$$

As households maximize their utility, all of the above equations must equal 0. Combining (2.57) and (2.56) yields:

$$\varrho_t w_t = \chi L_t^{\phi}, \tag{2.61}$$

where $\varrho$ is the marginal utility of consumption and is equal to $\lambda_t$ in (2.56):

$$\varrho_t = \frac{1}{C_t - hC_{t-1}} - \frac{\beta h}{C_{t+1} - hC_t}. \tag{2.62}$$

Inserting (2.56) in (2.60) yields:

$$1 = \beta\Lambda_{t,t+1}(1 + r_t^B), \tag{2.63}$$

where $\Lambda_{t,t+1}$ is the expected relative change in the marginal utility of consumption:

$$\Lambda_{t,t+1} = \frac{\varrho_{t+1}}{\varrho_t}. \tag{2.64}$$

Similar to (2.63), we derive the following equation for (2.58):

$$1 = \beta\Lambda_{t,t+1}(1 + r_t^D)\left(\psi_t - \Omega_D\left(\frac{D_t}{F_t^*}\right)^{\Omega_D}\right) + \frac{\Upsilon}{\varrho_t}(D_t + CBDC_t)^{\Gamma}, \tag{2.65}$$

and the following equation for (2.59):

$$1 = \beta\Lambda_{t,t+1}(1 + r_t^{CBDC}) + \frac{\Upsilon}{\varrho_t}(D_t + CBDC_t)^{\Gamma}. \tag{2.66}$$

To analyze the impact of the interest rate spread between $r^B$ and $r^{CBDC}$, we equate (2.59) and (2.60):

$$\beta \varrho_{t+1}(r_t^B - r_t^{CBDC}) = \Upsilon(D_t + CBDC_t)^{\Gamma}. \tag{2.67}$$

In equilibrium, the discounted real interest rate spread multiplied with the next period's expected marginal utility of consumption equals the marginal utility gained from holding liquidity. Since $\Gamma$ is negative, a decreasing interest rate spread will be offset by higher CBDC holdings — assuming that bank deposits are constant. Intuitively, a lower spread implies that households will keep more of their savings in the form of a liquid means of payment. Then, households do not consider the slightly higher interest income from bonds and the resulting additional consumption in period $t+1$ as worth giving up liquidity.

Equating the first-order conditions for CBDC (2.59) and deposits (2.58) yields:

$$\left[\frac{(1 - \frac{1+r_t^{CBDC}}{1+r_t^D} - \frac{\bar{N}-N_t}{\bar{N}}\Omega_N)}{1+\Omega_D}\right]^{\frac{1}{\Omega_D}} = \frac{D_t}{F_t^*}. \tag{2.68}$$

Note that the effect of liquidity is cancelled out, as deposits and CBDC provide the same liquidity services. The share of deposits to the total maximum external refinancing of banks $D/F^*$ depends on the interest rate spread between CBDC and deposits, the financial stress in the market, and the elasticity of the discount factor to changes in bank deposits $\Omega_D$. Note that, in the steady state, equality of interest rates implies that deposits are reduced to zero unless $\Omega_D$ reaches infinity. Intuitively, $\Omega_D$ determines households' subjective discount factor on bank deposits. Higher values of $\Omega_D$ 'push' $D$ closer to $F^*$ and, at the same time, reduce the interest rate elasticity of deposits.

The model cannot be solved as soon as we allow for the economically unreasonable case $r^{CBDC} \geq r^D$. First, there is no incentive for households to hold any deposits, thereby leading to negative values that imply a central bank refinancing over the maximum $F^*$. Second, a first-order approximation is not capable of capturing this non-linearity and produces misleading results. Therefore, we assume that $r^{CBDC}$ imposes a lower bound on $r^D$.

To compare bank deposits and government bonds, we equate (2.60) and (2.58):

$$\beta \varrho_{t+1}(1+r_t^B) = \beta \varrho_{t+1}(1+r_t^D)\left(\psi_t - \Omega_D \left(\frac{D_t}{F_t^*}\right)^{\Omega_D}\right) + \Upsilon(D_t + CBDC_t)^{\Gamma}. \tag{2.69}$$

In equilibrium, the discounted marginal utility gain from future consumption financed by interest income on bonds equals the same marginal utility from interest income on deposits, thereby accounting for subjective risk and the marginal utility from liquidity services.

To sum up, households' decision to allocate their savings depends on three dimensions: remuneration, liquidity, and risk.

### 2.6.2 Model Comparison with Gertler & Karadi (2011)

Our baseline model is based on Gertler and Karadi (2011). We adapt their model (hereafter referred to as GK) to make the introduction of a CBDC possible. The aim is to create a framework (1) that allows for changes in the level of deposits based on financial conditions and households' preferences and (2) that — before the introduction of a CBDC — preserves the main implications of Gertler and Karadi (2011) — that is, we retain the financial accelerator mechanism. This section outlines the implications of our implemented changes in households' maximization problem for the model output.

We make the following four assumptions. First, households actively choose between different forms of saving, accounting for differences in remuneration, liquidity, and risk. Second, banks do not merely intermediate funds from households to the production sector. Instead, they can additionally refinance themselves through the central bank. Third, the central bank fully allocates demanded funds to banks (full allotment) as long as their participation constraint holds. Fourth, refinancing via central bank money is more expensive than refinancing via deposits (Bindseil, 2020). These assumptions imply that an increase in central bank funds will offset a decline in households' deposits in the case of full allotment. Therefore, changes in deposits have only a minimal impact on total intermediated funds, capital, and production.

**Figure 2.7:** Baseline vs. Gertler & Karadi (2011)

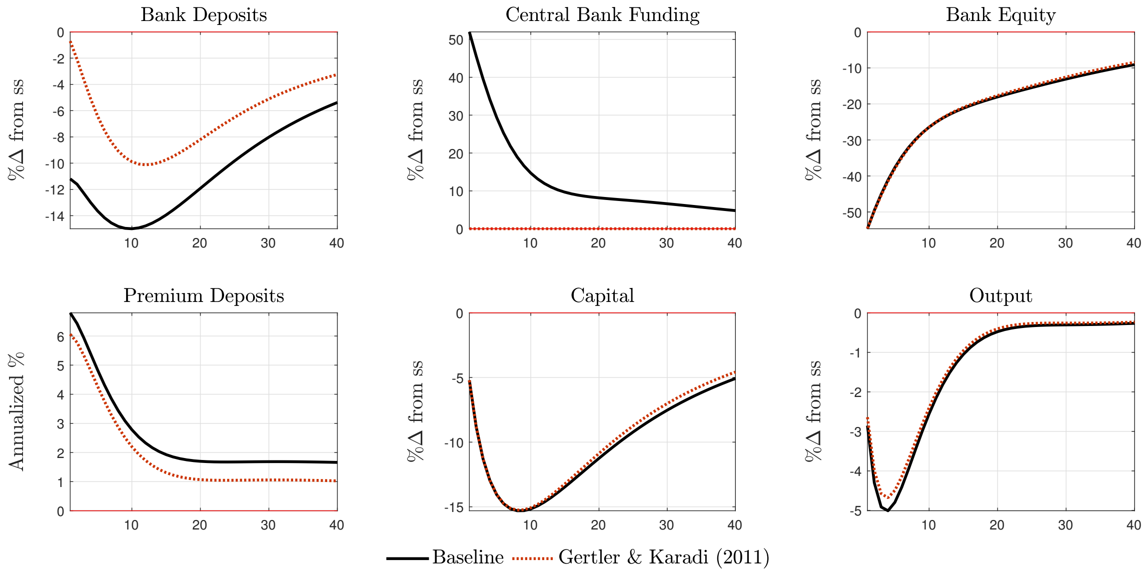


Figure 2.7 compares our model with GK. For both models, we induce a quality of capital shock of 5% with persistence 0.66 to simulate a crisis similar to the global financial crisis starting in 2007 (Gertler and Karadi, 2011). The fall in the quality of capital reduces effective capital and production. This reduction in production causes losses for intermediate goods producers and loan defaults. Hence, the losses are captured in a major decline in banks' equity — in our case, approximately 55%. Consequently, banks' participation constraint tightens, and households reduce their deposits. This reduction is amplified in our model, as households assign a risk to their deposits and distrust banks. As a result, banks have to substitute deposits with central bank funds. While the structure of bank funding is different for the two models, banks receive the same amount of total external refinancing, i.e., the roughly 10% difference in bank deposits between the models is offset by a 50% increase in central bank funding in our model. Nonetheless, driven by the loss in equity, total external refinancing and total intermediated funds decline over the following periods in both models and lead to a further reduction in capital and output — the financial accelerator effect. Less capital implies higher marginal productivity and grants banks higher returns. In combination with a decrease in the deposit interest rate, these returns yield higher premia on deposits. Consequently, banks quickly rebuild parts of their lost equity. However, with a declining premium,

this process slows down after 10 quarters and impedes further recovery processes. As a result, capital and output for both models remain below their steady states even after 40 quarters (10 years).

To sum up, our model — in contrast to Gertler and Karadi (2011) — allows for an active deposit decision of households, includes central bank refinancing, and features three different interest rates. Nevertheless, the model produces results similar to those obtained by Gertler and Karadi (2011) and retains their financial accelerator effect. Assuming full allotment, changes in bank funding structure do not affect the economy's overall performance.

### 2.6.3 Additional IRFs

In the following section, we present the remaining IRFs for the exercises conducted above. Note that we do not provide them for the simulations in Chapter 2.6.2. In addition, we exclude a few variables that do not provide additional information or that can be directly derived from the presented figures. The authors can provide additional material upon request.

**Figure 2.8:** Additional IRFs baseline vs. non-interest-bearing CBDC (with ELB)

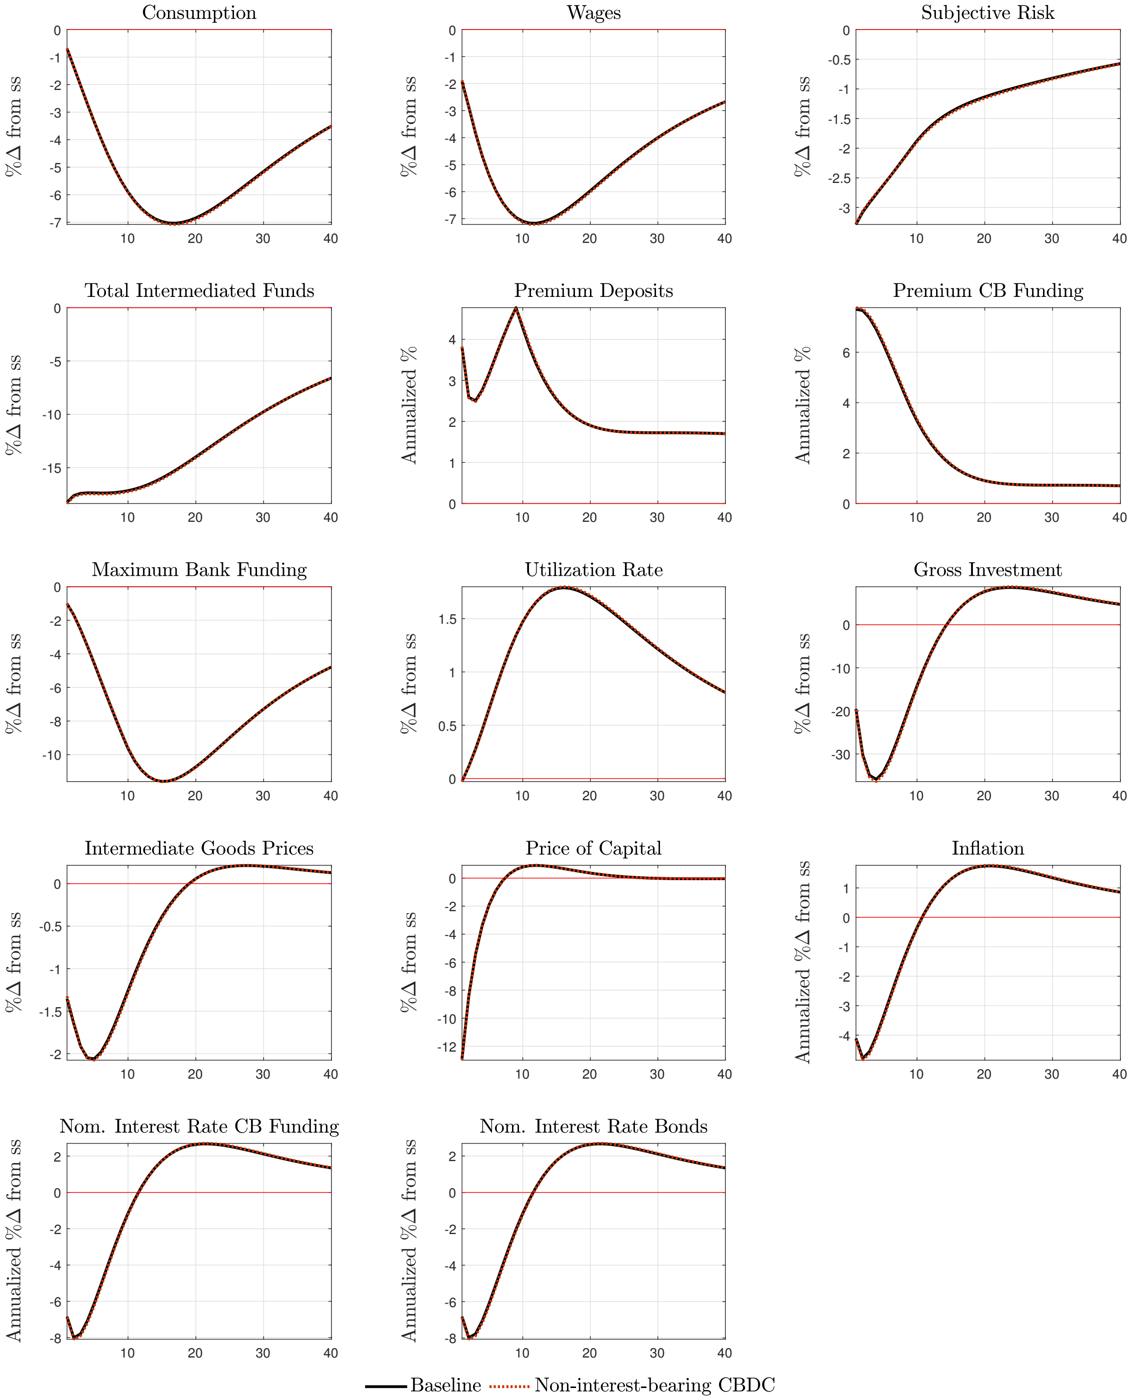

**Figure 2.9:** Additional IRFs baseline with ELB vs. interest-bearing CBDC

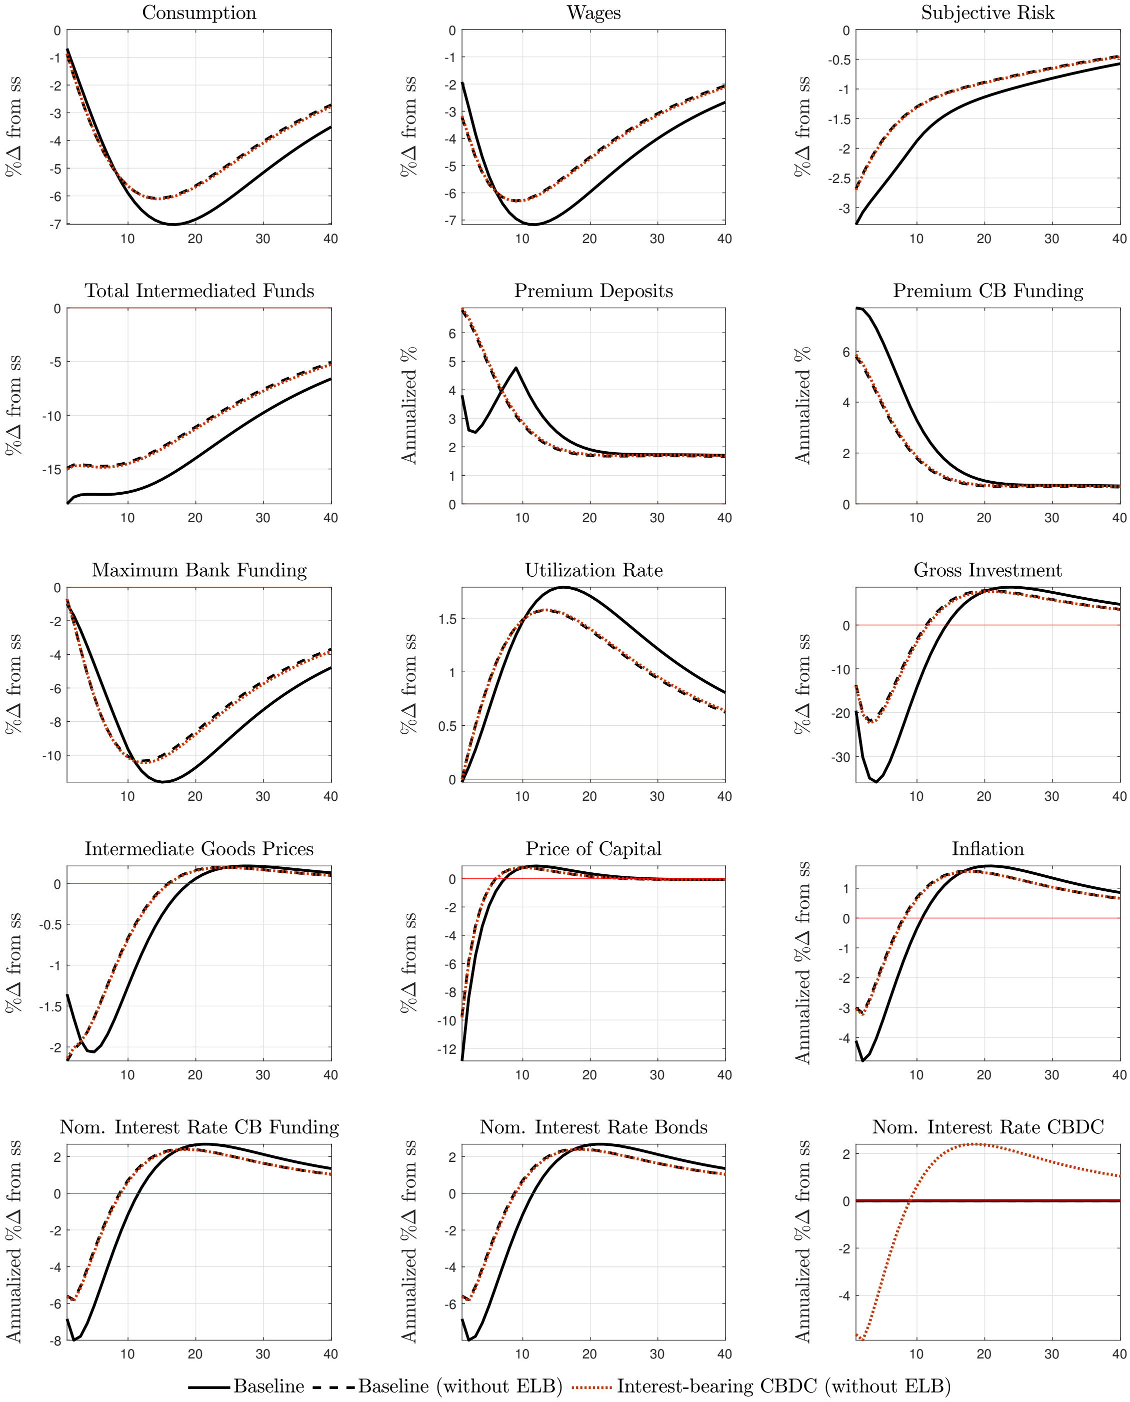

**Figure 2.10:** Additional IRFs interest-bearing CBDC with different allotment of central bank funds

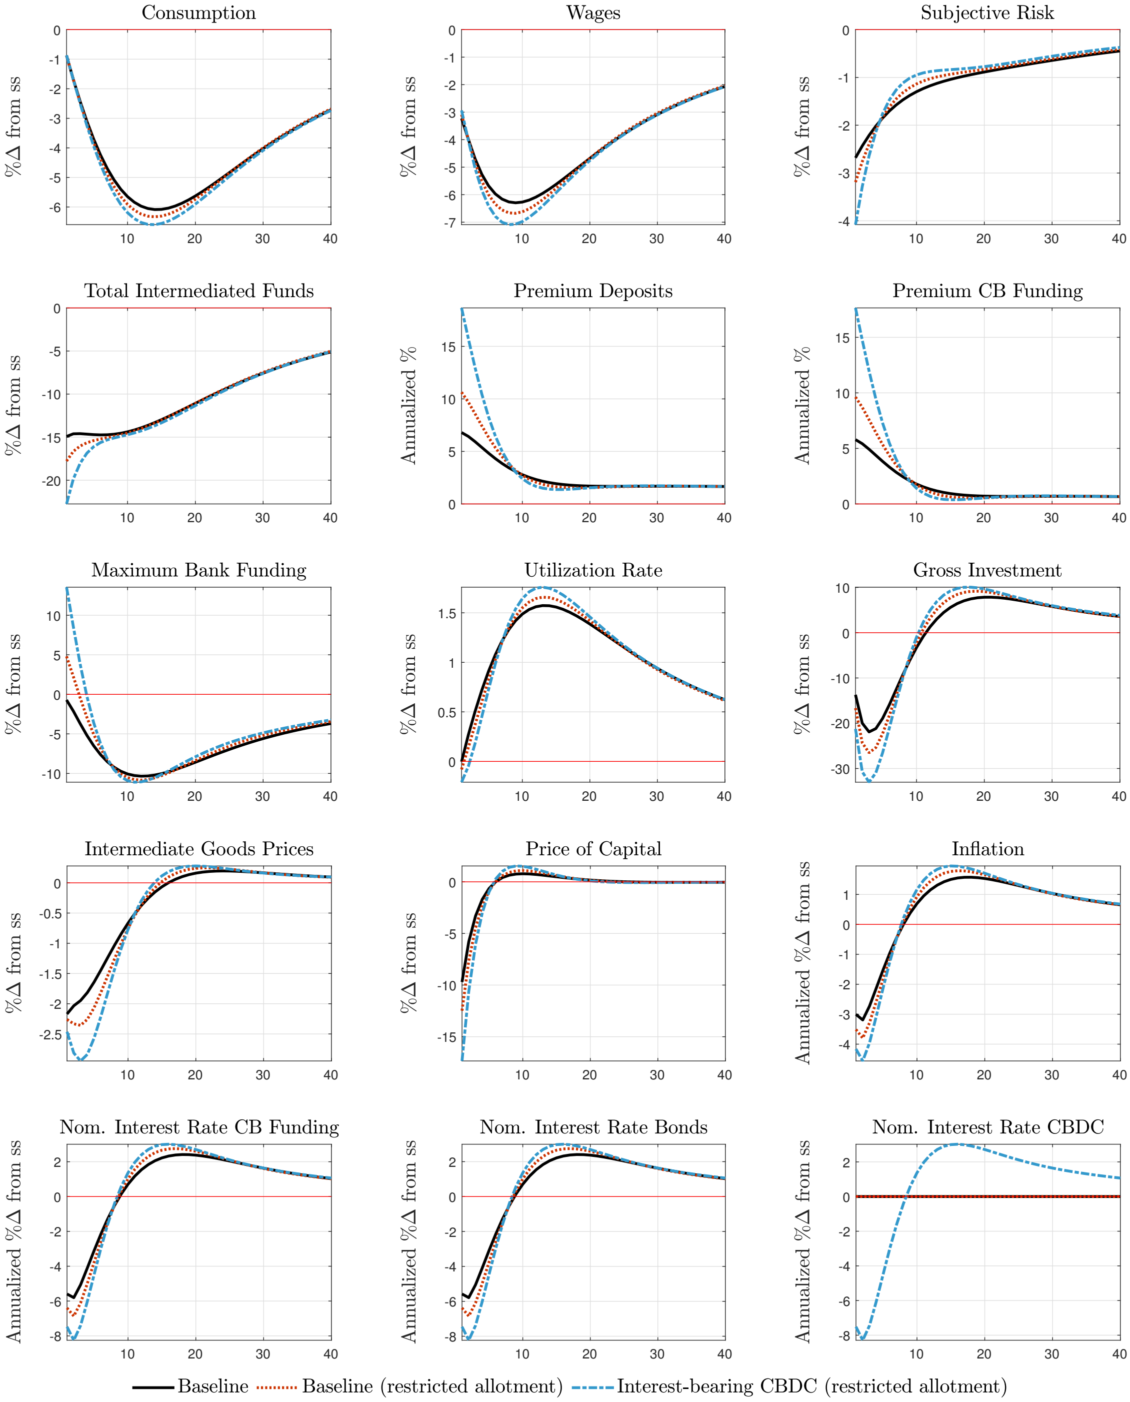

**Figure 2.11:** Flexible interest rate spread on CBDC with restricted allotment

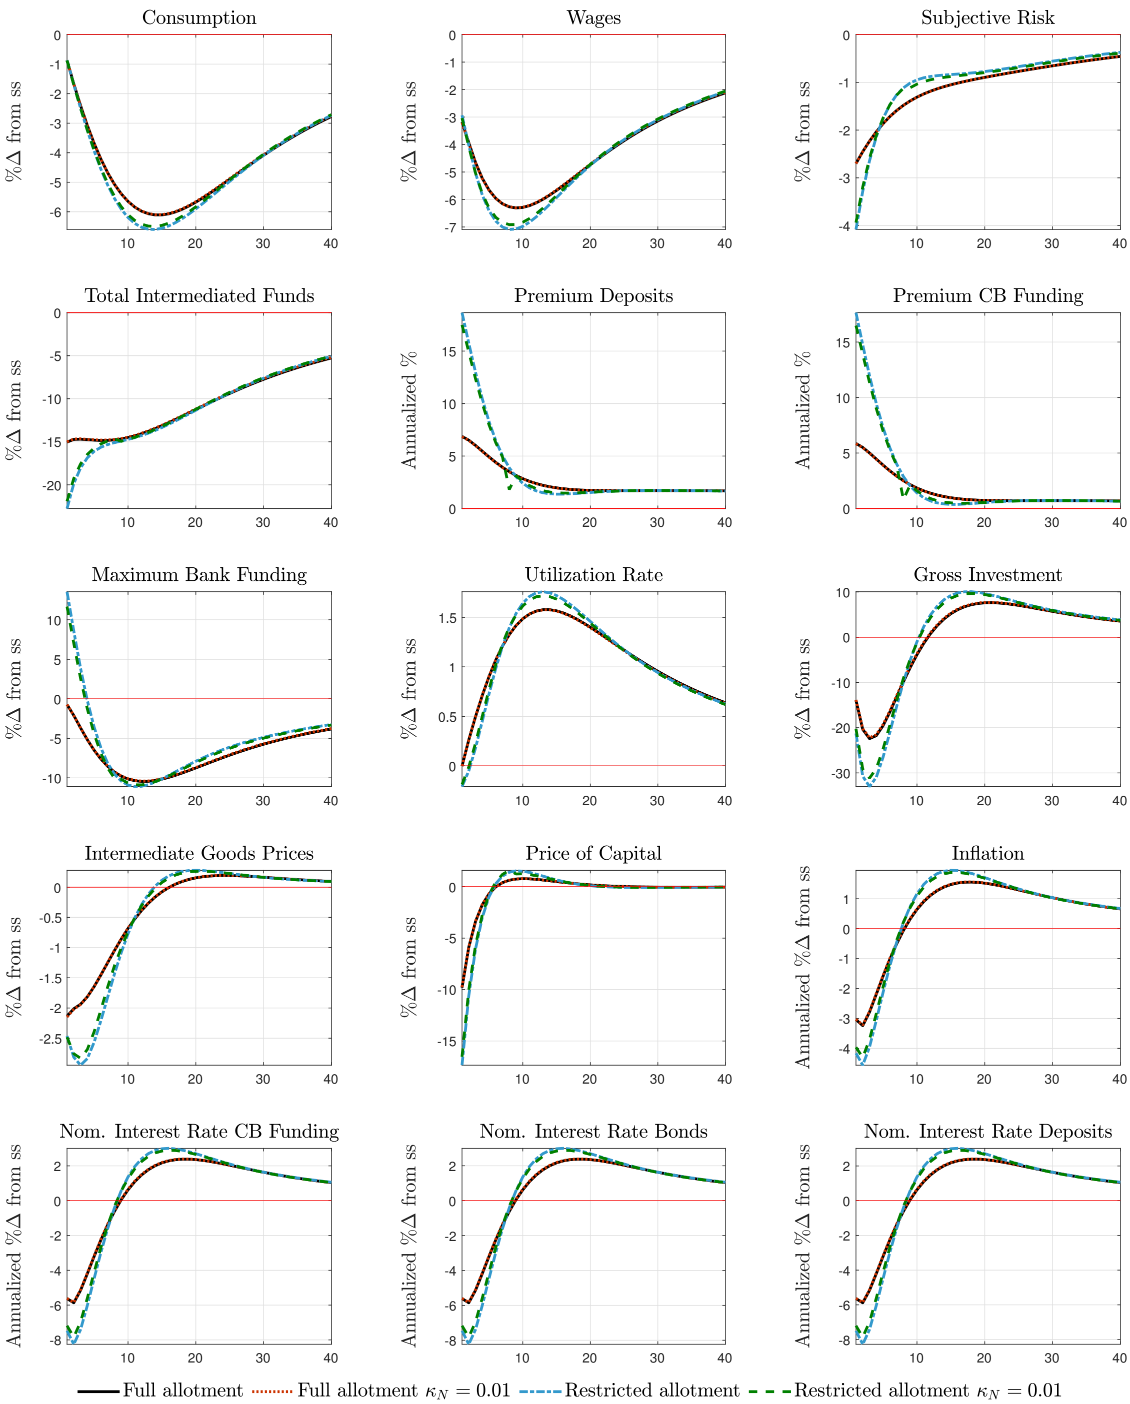

# Chapter 3

# Designing a Central Bank Digital Currency with Support for Cash-Like Privacy

**Abstract**

Most central banks in advanced economies consider issuing central bank digital currencies (CBDCs) to address the declining use of cash as a means of payment and to position themselves against increased competition from Big Tech companies, cryptocurrencies, and stablecoins. One crucial design dimension of a CBDC is the degree of transaction privacy. Existing solutions are either prone to security concerns or do not provide full (cash-like) privacy. Moreover, it is often argued that a fully private payment system and, in particular, anonymous transactions cannot comply with anti-money laundering (AML) and countering the financing of terrorism (CFT) regulation. In this paper, we follow a design science research approach (DSR) to develop and evaluate a holistic software-based CBDC system that supports fully private transactions and addresses regulatory constraints. To this end, we employ zero-knowledge proofs (ZKPs) to impose and enforce limits on fully private payments. Thereby, we are able to address regulatory constraints without disclosing any transaction details to third parties. We evaluate our artifact through interviews with leading economic, legal, and technical experts and find that a regulatorily compliant CBDC system based on ZKPs that supports full (cash-like) privacy is feasible.

*Keywords:* CBDC, compliance, design science, privacy, regulation, zero-knowledge proof.

*JEL classification:* E42, E52, E58.

Chapter 3 is a discussion paper. An earlier version of this work has been published as Gross et al. (2021).

## 3.1 Introduction

The monetary system is changing. In many advanced economies, the use of cash as a means of payment has declined steadily over the last decade (European Central Bank, 2020b) and in an accelerated way during the COVID-19 pandemic. Moreover, public money faces increasing competition from novel, private sector-issued forms of money, such as cryptocurrencies and stablecoins, and from Big Tech payment systems (European Central Bank, 2020a). Consequently, central banks take actions to preserve their monetary sovereignty. In January 2021, 86% of central banks around the world considered issuing their own digital currencies, i.e., central bank digital currencies (CBDCs) (Boar and Wehrli, 2021). While the Bahamas have already launched a CBDC and some other countries have introduced CBDC pilots (e.g., China), most countries are still debating and analyzing design options. In this context, the appropriate degree of transaction privacy receives great attention. For instance, in its announcement to start a project on the digital euro, the European Central Bank (ECB) stressed that a two-year investigation phase aims to identify "the design options to ensure privacy and avoid risks for euro area citizens, intermediaries and the overall economy" (European Central Bank, 2021b).

Privacy of transaction data is crucial, amongst other reasons, to avoid identity theft, threats to personal security, data exploitation, and harassment based on potentially embarrassing but legal purchases (e.g., Choi et al., 2021; Kahn et al., 2005; Kahn, 2018; Chaum et al., 2021). Privacy is also considered essential from an economic perspective, as it can help to avoid price discrimination (Acquisti et al., 2016; Odlyzko, 2004), making privacy a public good (Garratt and Van Oordt, 2021). Moreover, privacy constitutes a fundamental civil right enshrined in Article 12 of the United Nations Declaration of Human Rights (United Nations, 1948), in Article 8 of the Convention for the Protection of Human Rights and Fundamental Freedoms (European Court of Human Rights, 1950) as well as in Article 7 and 8 of the Charter of Fundamental Rights of the European Union (European Convention, 2000). In the context of CBDC, a consultation of European citizens revealed that they see privacy as the *most* important requirement for a CBDC (European Central Bank, 2021a).

A CBDC that stores transaction details in a centralized database operated by the central bank or payment service providers (PSPs) on behalf of the central bank bears the risk of losing trust and causing security incidents such as data breaches, e.g., due to human misbehavior or cyber attacks. The hack of New Zealand's central bank in 2021 demonstrates that cyber risks are indeed a threat that should be taken seriously (The Guardian, 2021). Furthermore, if sensitive data are stored centrally, end-users have to *trust* the operator that the privacy promises will not be compromised in the future. However, in such setups, operators could potentially change their minds or secretly analyze (historic) transaction data and share it with further parties, thereby potentially undermining privacy and trust. Trust in a payment system that inevitably processes sensitive data can be increased by following a *privacy-by-design* approach in which customers do not need to trust the operator for privacy protection and where large-scale data breaches are naturally excluded. In this case, private data would only be stored with the end-users involved in a transaction and not aggregated in a centralized system, thereby providing *trustless privacy*. Today, cash is the only regulatorily compliant form of money that provides full privacy by design. As, in the euro area, cash constitutes legal tender, and merchants are obligated to accept cash payments, there is a legal guarantee for anonymous payments with cash. If payments are conducted digitally, e.g., through mobile payments, bank transfers, or credit cards, the transaction data is stored with the involved PSP. Contrary to public perception, cryptocurrencies such as Bitcoin and Ether do not ensure a high degree of privacy, as transaction details are stored on a public ledger, and the pseudonymous addresses that send and receive cryptocurrencies can often be traced back to users that control them (Biryukov et al., 2014) through taking metadata into account (such as IP addresses) and information from exchanges that need to conduct know-your-customer (KYC) measures (Silfversten et al., 2020). Against this background, privacy-oriented cryptocurrencies such as Zcash and Monero have been developed. They use cryptographic techniques such as zero-knowledge proofs (ZKPs) to enable fully private payments (Fauzi et al., 2019). However, these cryptocurrencies do not conform with prevailing regulations, as unlimited anonymous payments open the door for illicit activities, such as money laundering and terrorist financing (Silfversten et al., 2020).

To secure access to a fully private, regulatorily compliant form of money in an increasingly digital environment, a CBDC has to provide a high degree of transaction privacy and offer (at least) the same privacy-preserving features as cash. Multiple central banks have already indicated their willingness to consider privacy-enhancing features for their CBDCs (Lane, 2020; Bank of Canada, 2020; Panetta, 2021), and first CBDC solutions have been proposed by both central bankers and academic researchers that provide some degree of transaction privacy. Naturally, these suggestions also consider regulatory constraints. However, software-based CBDC designs proposed by central banks, e.g., the ECB's anonymity voucher proposal (European Central Bank, 2019b), or by academic researchers (Dold, 2019; Tinn and Dubach, 2021; Chaum et al., 2021), do not support fully private transactions. Besides software-based designs, CBDC solutions can use hardware elements, e.g., used in computers, mobile phones, or smart cards, as gateways to access the CBDC infrastructure (European Central Bank, 2020a) and could, therefore, technically replicate the trustless privacy guarantees of cash. As an example, a hardware-based instrument is being tested for the CBDC in the Bahamas (Mastercard, 2021). However, such hardware-based solutions still exhibit considerable security challenges (European Central Bank, 2021b; Chaum et al., 2021).[40] Chaum et al. (2021, p. 11f) argue that experience has shown that

> *"any economically producible device that stores tokens with monetary value in an individual's possession [...] will be the target of successful forgery attacks as soon as the economic value from an attack would be sufficiently large."*
>
> — Chaum et al. (2021, p. 11–12)

Mitigating the risks of sophisticated forgery and further attacks would likely compromise privacy guarantees. Fortunately, the maturity of privacy-enhancing cryptographic techniques, and particularly of ZKPs, has grown considerably in recent years, offering new opportunities for enhanced privacy. ZKPs have already seen considerable adoption in the context of privacy-oriented cryptocurrencies, where they

[40] We recommend reference to the comprehensive discussion in Chaum et al. (2021) for a more detailed overview of the challenges associated with hardware-based solutions, particularly if deployed on a larger scale, and to Grothoff and Dold (2021) for additional arguments why a software-based CBDC design may be beneficial.

are used to ensure the integrity of payment systems, e.g., to prevent double-spending, while maintaining a high degree of privacy for users. However, ZKPs can be more broadly employed, with particular emphasis on enforcing further monetary or regulatory rules in a privacy-oriented payment system. Literature on cryptography already acknowledges the suitability of ZKPs for reconciling privacy and integrity or compliance requirements for electronic payments, e.g., through imposing turnover or per-transaction limits (e.g. Garman et al., 2016). Bontekoe (2020) specifically proposed an extension of Zcash in which third parties escrow users' digital identities and ZKPs allow the enforcement of turnover limits.

Still, regulators, central bankers, and researchers have repeatedly claimed that reconciling full privacy with regulatory constraints is not possible (e.g. Auer and Boehme, 2021; Armelius et al., 2021). This statement seems to indicate a lack of communication between the different research streams. Moreover, neither CBDC nor cryptographic literature has so far provided a rigorous, holistic evaluation of a payment system design that addresses regulatory requirements while supporting fully private payments and that evaluates the design with key stakeholders, such as central bankers and regulators.

To address this research gap, we follow a design science research (DSR) approach to design and evaluate a holistic, software-based CBDC system that is based on ZKPs and supports fully private payments. We first consolidate proposals from the cryptographic and CBDC literature to develop an account-based CBDC payment system that is fully private by design while addressing regulatory requirements by using per-transaction, turnover, and balance limits. We also instantiate our design through an implementation of the core transaction types using ZKPs. We then evaluate and refine our IT artifact in four evaluation cycles consisting of a total of 22 interviews with 44 experts in the areas of regulation, cryptography, central banking, identity, and payments. We find using ZKPs for CBDCs can replicate cash-like privacy in the digital realm and ensure adherence to regulatory constraints. Against this background, ZKPs enable strict privacy protection by design, storing personal transaction data only on the end-users' devices (*trustless privacy*).

The theoretical contribution of our paper is twofold: First, our innovative software-based CBDC payment system combines elements from different strands of the literature, including cryptography, privacy-by-design concepts, and CBDCs. Second, the evaluation of our CBDC system by key stakeholders in the cadre of DSR allows us to assess the practical feasibility of a CBDC design that provides cash-like privacy using ZKPs and also to discuss risk mitigation measures.

Our paper is structured as follows: In Chapter 3.2, we introduce essential background knowledge on CBDCs, different notions of transaction privacy, regulatory aspects of CBDCs, and ZKPs. We then present our DSR approach in Chapter 3.3. Subsequently, we discuss related work and describe the design of our IT artifact in Chapter 3.4, followed by the presentation of our evaluation cycles in Chapter 3.5. Chapter 3.6 discusses general implications of our approach, including implications for monetary policy. Chapter 3.7 summarizes our main findings, describes limitations, and gives an outlook on future research opportunities.

## 3.2 Theoretical Foundations

### 3.2.1 Central Bank Digital Currencies

In general, there are two forms of CBDCs, wholesale and retail CBDCs (Bech and Garratt, 2017). A wholesale CBDC is a digital form of central bank money accessible for financial institutions to optimize the settlement of wholesale payments and tokenized financial assets. A retail CBDC, in contrast, constitutes a novel form of central bank money available to the general public. In this paper, we solely refer to a retail CBDC, as we focus on end-user payments. A retail CBDC unites features of today's predominant forms of money: cash and bank deposits (Bech and Garratt, 2017). While cash is issued by central banks in physical form, bank deposits are issued by commercial banks in digital form. As central bank money, CBDCs bear no counterparty risk because the central bank issuing CBDC cannot – by definition and in contrast to commercial banks – go bankrupt.[41] CBDCs would hence provide a safer and practically riskless form of money for end-users.

CBDCs can be designed and implemented in different ways (European Central Bank,

[41] Today, deposit insurance schemes are established to address the risk of commercial bank money and avoid that, in the case of bankruptcy of a commercial bank, customers face substantial financial losses. However, deposit insurance schemes are not available in all countries equally, and commercial bank money is only secured until a specific threshold, e.g., in Germany up to 100,000 euro per client per financial institution.

2020b; Kiff et al., 2020; Auer and Boehme, 2020). Auer and Boehme (2020) identify *architecture*, *access*, and *technology* as the three main design considerations for a CBDC.[42] The architecture model specifies the role of the central bank and other market participants in the CBDC ecosystem. The account management, onboarding processes, and distribution of a CBDC might be conducted directly by the central bank (direct model) or by private sector PSPs (intermediated model). The access model defines how CBDC transaction data is stored and how access is managed. In an account-based model, the CBDC is stored in accounts, and hence the ownership of a CBDC is tied to an identity. In a token-based model, the central bank issues digital bearer instruments and ties the CBDC ownership to the (proof of) ownership of the CBDC units itself, similar to cash today. Regarding technology, a CBDC can be issued either via a centralized or a distributed ledger. If a centralized ledger is used, the central bank manages and controls the CBDC system. In the case of a distributed ledger, data processing, storage, and governance can be distributed across additional private or public sector institutions.

### 3.2.2 Privacy and Regulatory Compliance of Payments

In this paper, we distinguish between private, anonymous, and fully (cash-like) private transactions. In a *private* transaction, the transaction amount remains unknown, but the sender and receiver, i.e., the transaction parties, might be known to third parties (e.g., PSPs, the central bank, or regulatory authorities). In an *anonymous* transaction, the identities of the sender and receiver remain hidden, but the transaction amount might be known. *Fully private* transactions are private and anonymous; neither the transaction amount nor the sender or receiver are revealed to third parties. Therefore, our definition of full privacy is similar to the concept of secrecy as the concealment of information (Tefft, 1980; Bok, 1989). Full privacy or secrecy describes the attempt of consumers to avoid sharing information in order to prevent third parties from creating a digital representation of the real self (Zwick and Dholakia, 2004; Dinev et al., 2013).

Today, fully private and regulatorily compliant transactions are only possible with

[42]The fourth dimension refers to retail and wholesale interlinkages. Such interlinkages are especially relevant for cross-border CBDC payments. As we abstract from cross-border use in this paper, we do not consider this dimension.

physical cash and, to a certain extent, with e-money. All other payment methods are either not fully private (e.g., credit cards, bank transfers, Bitcoin), or they are not regulatorily compliant (e.g., privacy-oriented cryptocurrencies such as Monero and Zcash). In order to restrict the large-scale financing of illicit activities, regulators usually enforce per-transaction, turnover, and/or balance limits for anonymous payments. For instance, there are *turnover limits* for fully private cash payments in many euro area countries such as Greece (500 euros), France and Portugal (1,000 euros), Italy (2,000 euros), Spain (2,500 euros), Belgium (3,000 euros), and Slovakia (15,000 euros) (Pocher and Veneris, 2021). For anonymous e-money transactions, the 5th Anti-Money Laundering Directive specifies a monthly *turnover* and *balance limit* of 150 euros (European Union, 2018). As, to date, regulatory frameworks do not capture CBDCs, there are no such regulatory limits for CBDCs yet. However, it seems reasonable to expect that similar limits would need to be introduced for anonymous CBDC payments, similar to today's restrictions for anonymous cash and anonymous e-money transactions.

### 3.2.3 Zero-Knowledge Proofs

The notion of ZKPs was first introduced in the 1980s, describing "proofs that convey no additional knowledge other than the correctness of the proposition in question" (Goldwasser et al., 1989, p. 186). ZKPs refer to cryptographic protocols in which a *prover* can convince a *verifier* about a mathematical statement, for example, that the prover knows a piece of data that has specific properties. This statement may refer to the knowledge of a pre-image of a publicly known value under a hash function or about properties of the result of a publicly known algorithm that was executed on public or private data. In this setting, with a ZKP, the prover can convince the verifier without disclosing *any* information beyond the statement under consideration (Ben-Sasson et al., 2013; Ben-Sasson et al., 2018). If the statement refers to the output of an algorithm, a ZKP can enforce *computational integrity* without the need for the verifier to replicate the computation. Besides providing *confidentiality* for data and intermediate steps in a computation, an appealing property of many ZKPs is that they are *succinct*, i.e., the size of proofs and the computational complexity required to verify them is significantly smaller than applying the algorithm. In the case of the zero-knowledge succinct non-interactive argument of knowledges (SNARKs) we use in our CBDC system, both proof size and verification complexity are even *independent* from the complexity of the computation that is to be verified (Ben-Sasson et al., 2013). However, general-purpose ZKPs that can cover a large class of statements come at a high computational overhead for the prover.

In the 25 years after their discovery, researchers have leveraged special types of ZKPs in some contexts, such as enforcing *correct behaviour* in multiparty computations or *selective disclosure* in digital identity management schemes with *anonymous credentials*. The latter describes digital certificates that a trusted organization signed digitally and that their owner can use to prove claims about parts of the content of these certificates without revealing all of the contained information. In particular, when verifiably presenting attributes attested in the anonymous credential, strongly correlating contents such as the value of the digital signature itself do not need to be revealed (Ben-Sasson et al., 2013). Lately, these anonymous credentials have seen first adoption in so-called decentralized or self-sovereign identity projects, as explored by the public and private sector in Canada and Germany, among others (Kubach and Sellung, 2021). However, practical applications remained rare, as, for general-purpose ZKPs beyond these very specific cases, the computational complexity for the prover was prohibitive. Also, seemingly, there was not a considerable need for deploying ZKPs because information systems (IS) were generally designed with a service provider that was trusted with respect to both integrity and confidentiality. However, this paradigm started to shift with the advent of Bitcoin and the decentralization as facilitated by blockchain technology. Building on blockchain technology, a new type of IS, decentralized applications, emerged that do not involve a third party that is trusted with respect to integrity by performing computations redundantly (Rossi et al., 2019). However, the replicated execution of operations on blockchains immediately leads to considerable challenges from a scalability and confidentiality perspective (Ben-Sasson et al., 2018; Kannengießer et al., 2020).

In this context, general-purpose ZKPs started to find applications in privacy-oriented cryptocurrencies such as Zcash or applications on Ethereum such as Tornado-Cash, building on prior academic work (Ben-Sasson et al., 2014) to provide a Bitcoin-like payment system with fully private transactions. In the last few years, additionally, the succinctness of proofs has been leveraged by various projects on the Ethereum blockchain and novel cryptocurrencies. In zk-rollups, an untrusted third party batches many operations and proves the correctness of the resulting state transition with a ZKP. Through their ability to solve privacy challenges in cryptocurrency and blockchain projects (Partala et al., 2020), ZKPs have received increased attention in academia and business, and have hence considerably matured in terms of performance and applicability. Consequently, IS building on blockchains and general-purpose ZKPs have already seen first adoption in industry consortia that leverage blockchain technology, e.g., in the context of medical supply chains (Mattke et al., 2019).

## 3.3 Method

This paper follows a DSR approach (March and Smith, 1995; Hevner et al., 2004; Peffers et al., 2007) to design, develop, and evaluate a CBDC system that provides full privacy while addressing regulary requirements related to anti-money laundering (AML) and combating the financing of terrorism (CFT). We structure our paper as proposed by Gregor and Hevner (2013). To ensure methodological rigor, we use the widely accepted DSR methodology proposed by Peffers et al. (2007). Thus, we apply the following six steps procedure to derive our IT artifact: (1) Problem identification, (2) objectives definition, (3) design and development, (4) demonstration, (5) evaluation, and (6) communication (see Figure 3.1).

**Figure 3.1:** Design science research approach

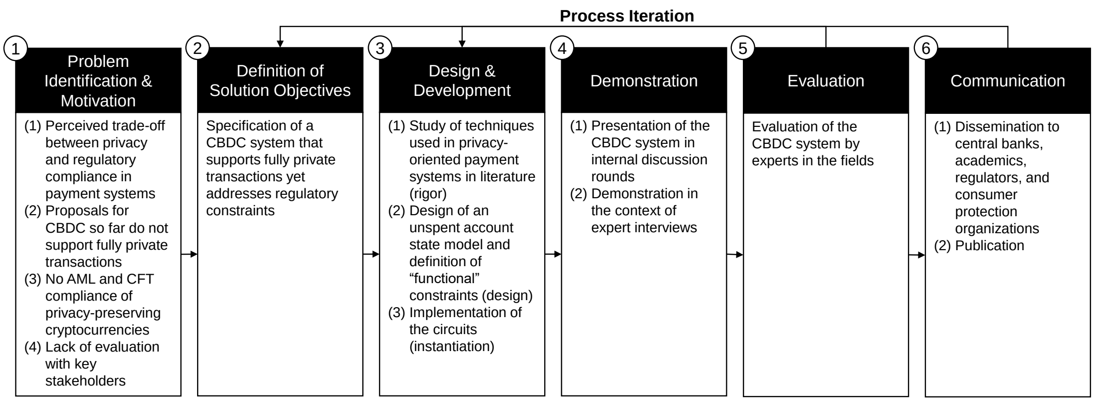

DSR was established to enable IS practitioners to find solutions to previously unsolved problems through a continuous build-and-evaluate process (March and Smith, 1995; Hevner et al., 2004). For an IT artifact to make a valuable contribution to IS research, it must address both a relevant business need (Hevner et al., 2004) and a general problem (Iivari, 2015). First, we identified the underlying problems and derived design requirements for our CBDC system.

We screened the most relevant primary literature on CBDCs, namely ECB (2020a) (European Central Bank, 2020a), BoC (2020) (Bank of Canada, 2020), Fed (2021) (Cheng et al., 2021), BoE (2020) (Bank of England, 2020), and BIS et al. (2020) (Bank for International Settlements et al., 2020), and identified both users' and central banks' requirements for a CBDC (see Figure 3.2).

**Figure 3.2:** Literature review on CBDC requirements

| | | **ECB (2020a)** | **BoC (2020)** | **Fed (2021)** | **BoE (2020)** | **BIS et al. (2020)** |
|---|---|---|---|---|---|---|
| Core | Liability of the central bank | ● | ● | ● | ● | ● |
| | Market neutrality | ● | ● | ● | ○ | ● |
| User req. | Security | ● | ● | ● | ● | ● |
| | Convenience and ease of use | ● | ● | ● | ● | ● |
| | Transaction speed and fast settlement | ● | ● | ● | ● | ● |
| | **Privacy protection** | ● | ● | ○ | ● | ● |
| | Usability | ● | ● | ● | ○ | ● |
| | Low cost | ● | ● | ○ | ● | ● |
| Central bank req. | **Regulatory compliance** | ● | ● | ● | ● | ● |
| | Resilience | ● | ● | ● | ● | ● |
| | Cooperation with market participants | ● | ● | ● | ● | ○ |
| | Universal access | ○ | ● | ○ | ● | ● |
| | Cost efficiency | ● | ● | ● | ● | ○ |
| | Interoperability | ● | ● | ○ | ● | ● |

*Note*: We included requirements that have been referred to in at least four of the five studies of central banks. If the dot is filled black (white), it indicates that the requirement is (not) mentioned.

We derived the following key requirements for end-users: privacy protection, high security, transaction speed and fast settlement, low costs, high usability, and availability. For central banks, the following requirements are important: AML and CFT compliance, market neutrality, resilience, cooperation with market participants, universal access, cost efficiency, and interoperability. In Chapter 3.1, we argued that privacy features should be at the core of a CBDC system, but that also regulatory constraints need to be addressed. In our DSR approach, we focused on these two core requirements. It is not (yet) feasible to address all requirements simultaneously in one artifact, as CBDC implementations are currently still at an early stage and, as there are trade-offs between different design parameters, such as resilience and fast settlement. Insights from the conducted expert interviews supported this hypothesis. The already proposed CBDC approaches do not enable fully private transactions, and related work in cryptography lacks a concrete design that can be used for the rigorous evaluation of their regulatory compliance and feasibility from the perspective of stakeholders (see Chapter 3.4.1). In this light, we proposed and instantiated a solution that uses cryptographic techniques, i.e., ZKPs, to address these requirements and to enable a discussion with stakeholders.

In particular, we aimed to develop a CBDC system that ensures cash-like privacy by design, where transaction amount and the identities of involved transaction parties are not shared with any third party. Compliance with regulation is enforced by per-transaction, balance, and turnover limits for fully private payments.

Next, we designed and developed our CBDC system in cycles that iterated between conceptualization, instantiation, and internal evaluation. We present the overall CBDC system architecture, onboarding procedure, depositing, withdrawing, and fully private transaction processes in Chapter 3.4. We then discussed our CBDC proposal in internal discussion rounds and presented it to leading experts from various fields. An overview of the interviewed experts is depicted in Figure 3.3. As one key result, our evaluation confirmed the feasibility and adequacy of our CBDC system. The adjustments to our CBDC architecture after each evaluation cycle are discussed in Chapter 3.5. As a final step, we disseminated our key findings to the interviewees and other stakeholders, in particular, to decision-makers in central banks and regulatory authorities, and to researchers. In addition, we published the source code of our prototype for the proposed CBDC system on GitHub.[43]

[43]The repository can be accessed at: https://github.com/applied-crypto/cbdc.

**Figure 3.3:** Overview of interviewed experts

| Cycle | Field of expertise | Int. No. | Exp. No. | Role | Organization |
|---|---|---|---|---|---|
| 1 | Law | 01 | 01 | Assistant Professor | University |
| | Law | 02 | 02 | Specialist Payment Fraud | Europol |
| | CBDC | 03 | 03 | Senior Financial Sector Expert | ex-IMF |
| | Payments | 04 | 04 | Head of Payments | Central Bank |
| | Payments | | 05 | Head of Digitalization and Payment Systems | |
| | CBDC | 05 | 06 | Chief Economic Advisor | SFB Technologies |
| | Cryptography | 06 | 07 | Global Managing Director Digital Assets | Accenture |
| 2 | IT | 07 | 08 | Technical Lead Digital Currencies | CBDC-developing company |
| | Economics | | 09 | Business Lead Digital Currencies | |
| | Business | | 10 | Product Manager Digital Currencies | |
| | Computer Science | | 11 | Data Engineer Digital Currencies | |
| | Information Systems | 08 | 12 | Senior Researcher | Research Institute |
| | Law | 09 | 13 | Banking Supervision Expert | Banking Association |
| | Payments | | 14 | Lead Digitalisation | |
| | Economics | | 15 | Chief Economist | |
| | Law | | 16 | Legal Lead | |
| | CBDC | | 17 | CBDC Expert | |
| | Payments | | 18 | CBDC Expert | |
| | Law | 10 | 19 | Professor | University |
| | IT / Business | 11 | 20 | Head of DLT Product | Bank |
| 3 | Economics | 12 | 21 | Alternate Member of the Governing Board | Central Bank |
| | Computer Science | | 22 | Professor | University |
| | CBDC | 13 | 23 | Senior Economist | International Organization |
| | Payments | | 24 | Senior Financial Market Analyst | |
| | Cryptography | 14 | 25 | Professor | University |
| | Economics | | 26 | PhD Candidate | Télécom Paris |
| | CBDC | 15 | 27 | Head of Blockchain | Association |
| | Payments | 16 | 28 | Market Infrastructure Specialist | Central Bank |
| | Computer Science | | 29 | IT Application Development Specialist | |
| 4 | Digital Identities | 17 | 30 | Head of SSI Consortium (IDUnion) | Main Incubator |
| | Law | 18 | 31 | Expert on CBDC and AML Regulation | University |
| | Business | 19 | 32 | Senior Manager Marketing & Public Affairs | Federal Money Printer |
| | CBDC | | 33 | Senior Project Manager CBDC | |
| | Digital Identities | | 34 | Senior Consultant Trusted Services | |
| | Cyptography | | 35 | Senior Principal Security Systems | |
| | Computer Science | | 36 | Technological Expert CBDC | |
| | Business | | 37 | Senior Account Manager | |
| | Business | | 38 | Regional Sales Director | |
| | Business | | 39 | Senior Business Development Manager | |
| | Computer Science | 20 | 40 | CBDC Technology Expert | Central Bank |
| | Computer Science | 21 | 41 | Research Group Lead | Research Institute |
| | Economics | | 42 | PhD Candidate | |
| | Finance | 22 | 43 | Research Group Lead | Consumer Protection Org. |
| | Economics | | 44 | Advisor | |

## 3.4 Our CBDC Design

### 3.4.1 Related Work

To conceptualize our design, we investigated research on CBDC as well as academic literature on cryptography with a focus on privacy-oriented digital payment systems. Chaum (1983) founded the research stream on *e-cash* that aims to develop cryptography-based payment systems that are private by design and make payments untraceable. He proposes a design in which users need to exchange their received *digital banknotes* for new ones in a compulsory interaction with a trusted PSP, e.g., a bank. To make the spending and receiving of a specific banknote unlinkable, *blind signatures* hide the serial numbers of unique and thus distinguishable digital banknotes. However, copying and thus double-spending these digital banknotes cannot be prevented technically. Instead, to hold users accountable, the cryptographic protocol allows retrieving a user's identity that is hidden in the digital banknote from combining the information obtained in two different payments with the same digital banknote. Despite being computationally very efficient, this approach implies that, while the sender remains anonymous, the receiver is identified by its PSP, and the payment amount is transparent. Moreover, the design cannot practically enforce per-transaction limits. Turnover and balance limits are cumbersome to implement because in a CBDC system that involves multiple PSPs, as currently planned in most jurisdictions, a synchronization among PSPs would be required to detect double-spending attempts and to prevent users from visiting multiple PSPs to circumvent these limits. Furthermore, it is not clear how to implement basic programmability functionalities, such as interest payments, on top of this design.
Sander and Ta-Shma (1999) address some limitations of Chaum's approach, for instance, hiding the transaction amount by dividing it into discrete shares. Nevertheless, the PSP still learns the transaction amount paid and received in this epoch as well as the identity of the receiver. The first e-cash system that hides the identity of both sender and receiver and also the transaction amount without the need for a trusted third party was proposed by Camenisch et al. (2006). In addition to hiding the identity of the receiver and the transaction amount, this payment system guarantees the sender's anonymity as long as they do not double-spend and

exceed a pre-defined per-transaction limit. Technically, the proposal is based on ZKPs that are mathematically tailored specifically to this use case. However, as this approach is still designed for digital banknotes and different payments in which the same individual is involved and cannot be identified, turnover limits cannot be enforced. Additionally, it is challenging to extend this model to incorporate basic programmability features.

While none of the previous academic work seems to have received considerable adoption, the first practical emergence of privacy-oriented digital payment systems happened in the context of cryptocurrencies. Researchers developed privacy-oriented modifications of Bitcoin (Nakamoto, 2008), as the use of a public distributed ledger increased the need for a privacy-enhancing design. The two arguably most relevant academic studies in this context are Zerocoin, which hides transaction parties but not the transaction amount (Miers et al., 2013), and Zerocash (Ben-Sasson et al., 2014), which additionally hides the transaction amount and laid the foundation for the well-known cryptocurrency Zcash. Both approaches use general-purpose ZKPs, in particular SNARKs, to demonstrate that transactions are valid and respect agreed-upon monetary policies (e.g., no double-spends) without revealing transaction details. Garman et al. (2016) acknowledge that "from an investigative standpoint, Zerocash is no different than cash" (Garman et al., 2016, p. 81). However, they also point to the conflict between regulation and private payment systems, as both Zerocoin and Zerocash do not take legal frameworks into account that restrict the use of anonymous payments to comply with AML and CFT regulation. Consequently, Garman et al. (2016) sketch potential extensions of the Zerocash model to address regulatory constraints, such as considering specific jurisdictions involved in payments, per-transaction limits, the tracing of specific coins, and the payment of appropriate taxes. Since exceeding the limits enforced by ZKPs would imply that no transactions are possible and thus reduce the utility of the payment system substantially, they propose a tiered approach for a private payment system so that any transaction above the spending limit should be additionally signed by an authority that conducts KYC and AML checks. Similarly, Bontekoe (2020) follows the approach of modifying Zcash by adding an account-based system based on a previous KYC-process. This setup allows limiting the transactions of an account within a specific epoch that can ensure balance and turnover limits and, thus, regulatory compliance. The study also describes the possibility of enabling transactions that exceed a limit by verifiably encrypting the associated transaction data and allowing for subsequent checks through a dedicated authority.

Both Garman et al. (2016) and Bontekoe (2020) also mention that users' digital identity management based on digital certificates can take a valuable role in this context. Literature that considers the role of identity in privacy-oriented payment systems also emphasizes that a mechanism for revoking identities and accounts is required (Choi et al., 2021). From a regulatory perspective, this resonates well with sanctions lists such as the Office of Foreign Assets Control's Specially Designated Nationals.

Alongside these developments, cryptographic research in the context of CBDCs continued pursuing Chaum's initial *asymmetric* approach to transaction privacy, offering anonymity for the sender but not for the receiver (Chaum et al., 2021; Tinn and Dubach, 2021). Chaum et al. (2021) base their design on blind signatures and hence employ similar techniques as earlier work by Chaum (1983) from a technical perspective, whereas the approach by Tinn and Dubach (2021) is based on SNARKs. Veneris et al. (2021) propose a system for CBDC that makes transaction amounts transparent yet enables anonymity through identity escrow (which is consequently not trustless). They also add a hardware-based solution that provides essentially cash-like privacy but requires regular online settlement. In these designs, the privacy protection of the sender covers many practical requirements, especially when a consumer purchases a product from a business that needs to disclose its accounting transparently. However, as discussed in Chapter 3.1, particularly for payments between end-users, trustless, cash-like privacy may be desirable.

To date, several CBDC designs have been proposed that aim for regulatory compliance and, at the same time, preserve users' privacy – at least to some extent (see Figure 3.4). However, these designs differ substantially in the extent to which privacy is guaranteed. In this context, *privacy by design* and *compliance by design* play a central role, representing systems where no trust in the operator is needed. Privacy by design and compliance by design can be achieved through extending a payment system like Zcash that allows fully private transactions with the possibility of person-related monthly turnover or transaction limits, as proposed by Garman et al. (2016) and Bontekoe (2020). Yet, since the perception that enabling fully private transactions in digital form fundamentally contradicts AML and CFT regulation is still widespread, and regulators, central bankers, and researchers have repeatedly claimed that reconciling full privacy with regulatory constraints is not possible (e.g. Auer and Boehme, 2021; Armelius et al., 2021), we present a detailed approach for a *privacy/compliance-by-design* CBDC. We also show how digital identities can replace an efficient and privacy-preserving KYC process.

**Figure 3.4:** Comparison of privacy-oriented CBDC solutions

| | | ECB (2019b) Anonymity Vouchers | Chaum et al. (2021) Blind signatures | Tinn and Dubach (2021) P-hybrid CBDC | Choi et al. (2021) Banking$^+$ and Cash$^+$ | Veneris et al. (2021) CBDL | Our approach UAS model with ZKPs |
|---|---|---|---|---|---|---|---|
| Privacy | Anonymity of the sender | ● | ● | ● | ◐ | ● | ● |
| | Anonymity of the receiver | ● | ○ | ○ | ○ | ● | ● |
| | Privacy of transaction amount | ○ | ◐ | ○ | ○ | ● | ● |
| | Trustless privacy (privacy by design) | ○ | ● | ● | ○ | ◐ | ● |
| Regulation | Regulation by design | ○ | ● | ○ | ○ | ◐ | ● |
| | Per-transaction limits | ● | ● | ○ | ○ | ● | ● |
| | Turnover and balance limits | ● | ● | ○ | ○ | ● | ● |
| | Anonymous onboarding | ○ | ○ | ○ | ◐ | ○ | ● |
| Other | Offline payments | ○ | ○ | ○ | ○ | ● | ○ |
| | Account-based system | ○ | ○ | ● | ◐ | ● | ● |
| | DLT-based system | ● | ○ | ● | ● | ● | ○ |
| | Involvement of PSPs | ● | ● | ◐ | ● | ● | ◐ |

*Note*: If the dot is filled black (white), it indicates that the attribute is (not) fulfilled. If the dot is half-filled the attribute is partially fulfilled.

In Figure 3.4, we compare our CBDC proposal with the existing privacy-preserving CBDC proposals. We mainly focus on aspects around privacy and regulation. While most of the CBDC proposals focus on asymmetric privacy, i.e., only the sender of money remains (fully) private (European Central Bank, 2019b; Chaum et al., 2021; Tinn and Dubach, 2021; Choi et al., 2021), our approach enables full (trustless) privacy by design for both the sender and the receiver. Our proposal also ensures regulation/compliance by design and utilizes per-transaction, turnover, and balance limits to restrict anonymous payments effectively. In contrast to the other approaches, in our system, users can be onboarded anonymously by proving the possession of a valid digital ID. However, we face the trade-off that a software-based CBDC providing high privacy guarantees comes at the cost of not enabling offline payments.

### 3.4.2 High-Level Overview of Our CBDC Architecture

Following Garman et al. (2016), our overall CBDC design is based on a two-tiered approach that supports both fully private and transparent CBDC payments. On the transparent side, the CBDC is distributed to end-users via PSPs – i.e., we use an intermediated model – and access is linked to an identity – i.e., we use an account-based model. Transaction data is stored in a centralized ledger. However, our system could also accommodate a distributed ledger. The disadvantage of a distributed ledger is that it may introduce challenges regarding performance, and privacy issues may arise for storing transaction- or account-related information on the transparent side. Thus, modifications, such as storing only PSPs' balances in an obfuscated way, may be necessary on the transparent side. In general, the transparent side is flexible to implement central bank-specific designs, e.g., introducing a maximum limit on transactions in general or using a direct instead of an intermediated model.

We propose an account-based model that allows the funneling of all transactions of a user into one single account, which has many similarities with the approaches by Garman et al. (2016) and Bontekoe (2020). Recently, the security of an account-based approach in combination with ZKPs has been studied in more depth also in Wüst et al. (2021), including several improvements in terms of performance. Yet, none of these publications incorporates a connection to digital identity management, which – as we will discuss later – is likely indispensable to address some of the challenges related to sharing access to the privacy pool. In our system, users maintain their accounting privately, and a transaction corresponds to updating their private account and sending a ZKP that proves to the central bank that the transaction's expected policies have been met. The transaction amount then essentially corresponds to the delta, i.e., the difference between the previous and the new account states' balances. In essence, users store and manage their own accounts and prove the correctness of their local accounting to the operator of the ledger, i.e., the central bank.

### 3.4.3 The Privacy Pool in Detail

Most cryptocurrencies, such as Bitcoin, record every transaction, including the sender, receiver, transaction amount, and authorization (in terms of the sender's digital signature) on a public ledger (Zhang et al., 2019). Additionally, transactions either point directly to one or more previously received but so far unspent coins (unspent transaction output (UTXO) model) or to the sender's and receiver's pseudonymous accounts with a public balance (account-based model). This setup generally allows one to track the transaction history of digital banknotes and corresponding metadata as well as additional information retrieved from exchanges to identify the involved parties in most transactions (Meiklejohn et al., 2016). Consequently, even if the ledger is not public, such a construction would allow the operator of the ledger to link transactions and correlate even pseudonymous accounts with real-world identities and, thus, to retrieve personal information.

In contrast, Zcash only stores cryptographic hashes of transactions in a ledger that hides details and ensures unlinkability. The payment details are only known to the sender and receiver. In particular, for every transaction (including its details), Zcash applies two different one-way cryptographic functions to generate two unique but different outputs: (1) a *commitment*, i.e., a hash consisting of the receiver address, the amount being sent, a number rho, and a random nonce, and (2) a *nullifier*, i.e., a hash consisting of the users' spending key and the secret rho from the commitment (Zcash Team, n.d.). The commitment and nullifier hence essentially both hide transaction details but are computationally infeasible to correlate without knowledge of the transaction details. To spend a previously received transaction in the form of a commitment, the corresponding nullifier is then published together with a ZKP that the nullifier corresponds to a previously published commitment (proof of knowledge of the joint pre-image) (Ben-Sasson et al., 2014). Then, the central bank has to check that the associated commitment has not been spent before by checking whether this particular nullifier has been published before. The ZKP-based construction hence certifies that the sender has access to digital banknotes that have not been spent previously and that the amount of money spent equals the amount received, without the need to disclose *any* additional information about the transaction, such as sender, receiver, amount, or a direct reference to the previous transaction in which the banknote was received (Ben-Sasson et al., 2014).

The design of our privacy pool is based on the construction of Zcash with an append-only ledger that stores commitments and nullifiers. Adding a per-transaction limit to Zcash would be an easy task (Garman et al., 2016). However, our approach contains a crucial modification, replacing the UTXO-based model with an unspent account state (UAS) model to (privately yet provably) funnel all of a users' transactions into one account and hence enable a user to prove that balance and turnover limits are also satisfied. Each commitment and nullifier thus represents a unique account state. By publishing a commitment and a new nullifier, the previous account state is invalidated and a new one is created. To validate transactions, the central bank receives the associated commitments, nullifiers, and proofs of the correct update of the account state using ZKPs. In particular, the central bank verifies the validity of the ZKPs and that the nullifier has not been revealed before and then adds the transaction to the ledger by including the commitment and nullifier in the existing associated ledger. Due to the succinctness property of the ZKPs used, the workload for the central bank in verifying the transactions is very small (Ben-Sasson et al., 2013). The pre-image resistance of hash functions (with a pre-image of high entropy) and the unlinkability of commitments and nullifiers ensure that information that is revealed to the central bank is not sensitive (Ben-Sasson et al., 2014). Consequently, only the account owners know their respective transaction and account details, such as the amount, current balance, and turnover. Nevertheless, they have to prove compliance with predefined rules by using ZKPs. This approach ensures *privacy by design* and enables the creation of *proofs of compliance with limits by default*. In such a system, entities that maintain the ledger only need to be trusted with respect to *integrity* and not with respect to protecting the users' privacy.

As common in the context of blockchains, for storing the commitments and enabling efficient proofs of inclusion for a commitment, we use Merkle trees. A Merkle tree is a cryptographic data structure that represents many entries by one identifier, i.e., the Merkle root (Merkle, 1987). This setup allows for proofs of inclusion that can be checked only by comparing with the Merkle root. The Merkle root changes with each new entry in the tree. By using these Merkle trees in combination with ZKPs, it is possible to prove that a transaction proposal refers to an existing commitment in the ledger without pointing to (and thus revealing) it. If the ledger is not public, as it is for a centralized ledger setup or a permissioned blockchain, it is necessary to prevent the correlatability of commitments and nullifiers that may occur through querying the Merkle path of a specific commitment, e.g., through querying Merkle subtrees.

To ensure the integrity of the payment system and the compliance with turnover, balance, and per-transaction limits, our UAS model contains the following information that is stored in digital wallets and consequently is only known to their respective holder:

- Identity information: A public key and a digital certificate that includes the account owner's identity information (digital ID card).
- Balance: The balance of the account holder.
- Epoch turnover: An accumulation of all amounts of spending transactions in the current epoch (whose length can be specified by the regulator).
- Epoch reset: The last reset of the epoch (the last time the epoch turnover was set to zero).

This structure should be seen as an initial proposal that can be flexibly extended when more account details need to be checked for compliance, e.g., one could include the nationality or the type of the account, i.e., private or corporate.

For our CBDC system, we determine three core transaction types for interacting with the privacy pool, illustrated in Figure 3.5. In the following, we outline and discuss these three types of transfers that are related to the privacy pool – onboarding, semi-private transfers, and fully private transfers. All three types of transactions involve private inputs, public outputs, and a ZKP that connects them and proves the correctness of the local accounting to the central bank. The account owner

**Figure 3.5:** Transaction types in the privacy pool

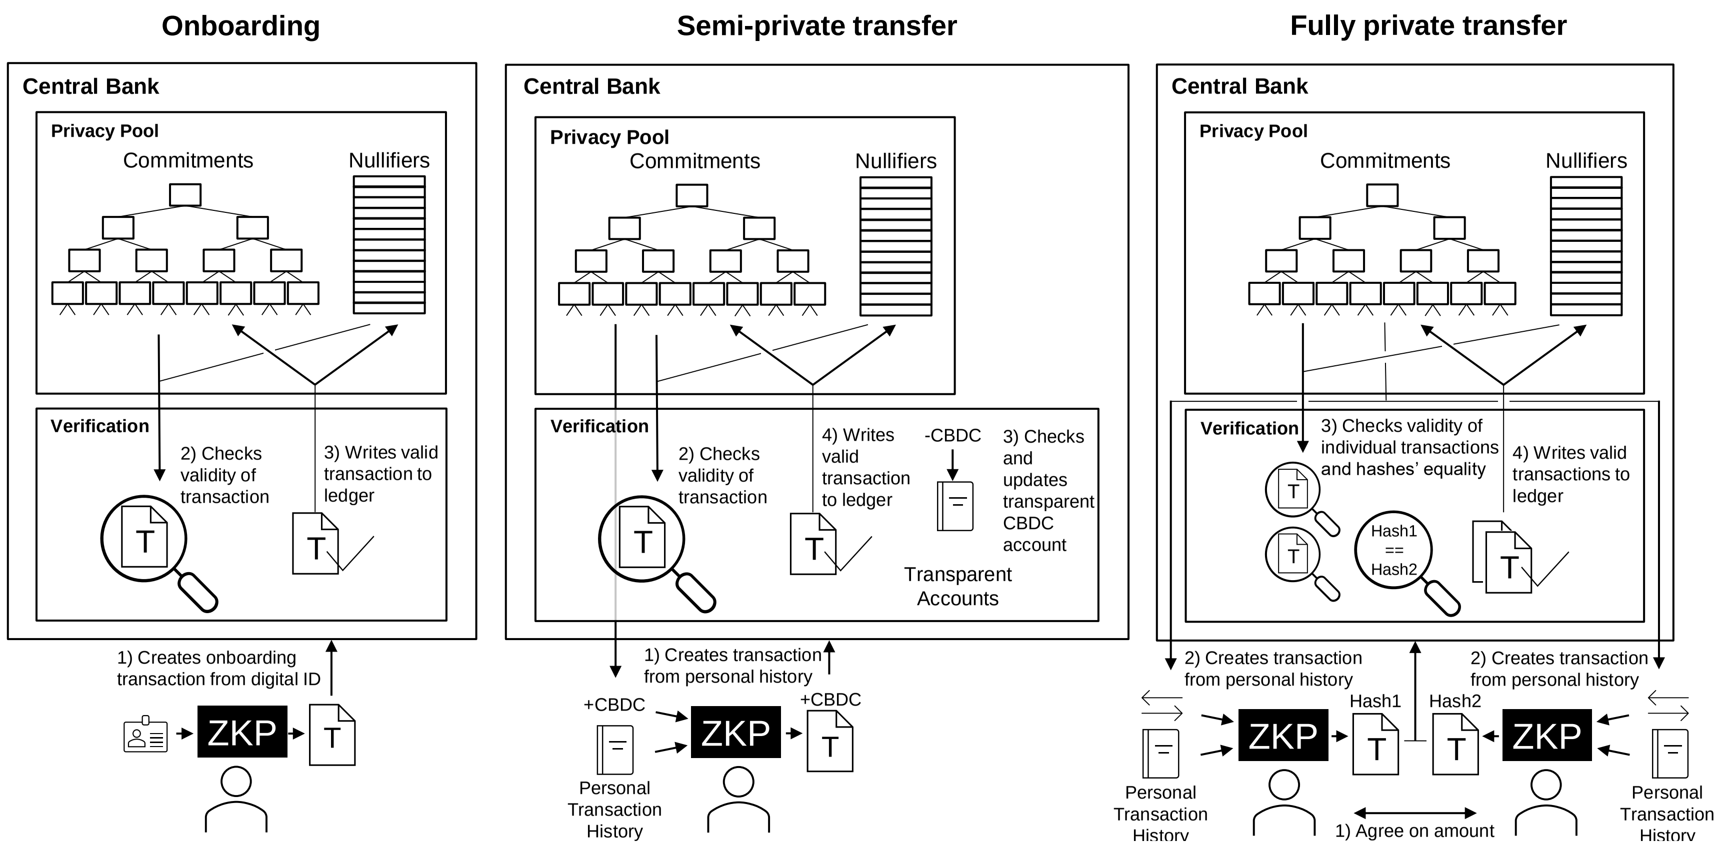

provides private inputs, which contain the already mentioned account information. This data is sensitive, therefore, it stays hidden. However, the private inputs result in the public outputs using, inter alia, one-way functions. Hence, the account owner shares the public output (commitment, nullifier, etc.) with the central bank along with a ZKP that ensures that the public outputs were computed from the private input according to the rules of the payment system. In the following, | denotes the concatenation operation.

### Onboarding and one-time registration

Each account is tied to an individual cryptographic key pair and, using a digital certificate, to a government-issued identity. We assume that each user owns a single digital ID card that is also bound to a cryptographic key pair. Many countries already integrate keypairs into their physical ID cards and even provide dedicated mobile apps to store the ID card in digital form, e.g., Germany and Estonia. Note that our CBDC proposal could also be implemented based on a centralized or decentralized database instead of a digital ID card. These databases would need to document which users have already registered for the privacy pool to ensure that every user can only create one account. The decentralized alternative requires substantial coordination between various PSPs about the registered participants, implying higher governance efforts and potential privacy compromises. Using a centralized database would imply higher incentives for users to circumvent the limits because not the ID itself but only the CBDC-storing device would need to be passed on.

To onboard a new user (one-time registration), the user has to create a cryptographic proof that they possess a valid ID that is not expired or revoked and that the initial commitment is deterministically derived from the ID card via a one-way function and contains the correct initial account entries. Therefore, each onboarding procedure of one individual is based on the same key pair and always results in the same commitment. In detail, a ZKP for onboarding and, thus, opening a new account would be structured as follows:

**Onboarding**
INPUT: Key pair and digital certificate (government-issued ID card), timestamp
OUTPUT: Commitment to initial account state, timestamp, ID issuer's public key
CHECKING THAT:

- *the account holder controls a valid ID card (signature, binding, non-revocation)*
- *the identity lodged in the account corresponds to the ID card*
- *the account's initial balance and epoch turnover are set to zero*
- *the account's initial epoch reset equals the timestamp*
- *the commitment equals the hash of the signed initial account state*

The onboarding process works as follows: First, using their legit key pair and digital ID, users create the onboarding transaction, consisting of the initial commitment and the ZKP, and send it to the central bank. Second, the central bank verifies the validity of the ZKP that refers to the commitment and adds the commitment to the current state of the ledger. Since the key is only used for the generation of the ZKP and the commitment represents the encryption of a hash, the central bank does not learn anything about the onboarded user. The central bank can also detect multiple attempts to create an onboarding transaction, as the commitment is deterministically derived from the ID card.

It is important to note that the anonymity of the onboarding process is *not* required to guarantee the anonymity of a user's subsequent transactions in the privacy pool due to the unlinkability of commitments and nullifiers. Consequently, the central bank could even demand that users that register for the privacy pool present some of their identity attributes, e.g., their nationality or the issuing authority of their digital ID.

Notably, it is not even necessary to check whether the commitment has already been used for previous onboarding to ensure that every user only has a single account: the commitment can only be spent once, independent of how often it is included in the Merkle tree because the corresponding nullifier is also unique. Nevertheless, to avoid attacks that spam the ledger with correct yet useless transactions, it may be useful to check for such collisions.

#### Semi-private transfers

Semi-private transfers describe the exchange of funds between an account in the privacy pool and a transparent account. These transfers include deposits and withdrawals from the same user so that a user can transfer CBDCs from their transparent account to their privacy pool account, or vice versa. In a semi-private transfer, a user combines an update of their account in the privacy pool with an update of their transparent account that is confirmed by the PSP. As it has to be ensured that the total money supply in the system is unchanged when money is deposited or withdrawn, the transaction amount, i.e., the difference in balances between the spent and created account state, must be disclosed to check whether it is equal to the counter-transaction on the transparent side.

In the following, we consider the depositing process as an example of a semi-private transfer. In more detail, the ZKP would be specified as follows:

> **Semi-private transfer**
> INPUT: Key pair, Merkle path of the previous commitment,
> previous account state, amount
> OUTPUT: Merkle root, nullifier, new commitment, (deposit/withdrawal) amount
> CHECKING THAT:
> - *the previous commitment is contained in the tree represented by the Merkle root*
> - *the previous account state belongs to the previous commitment*
> - *the nullifier equals the hash of the previous account state*
> - *the new account state is correct (e.g., new balance = old balance + amount)*
> - *the new account state complies with the rules*
>   *(e.g., positive balance, epoch turnover below turnover limit)*
> - *the commitment equals the hash of the signed new account state*

First, the user creates a new commitment and nullifier and attaches a ZKP proving that the account update is legitimate and corresponds to the public outputs. Second, the central bank verifies the ZKP and checks if the public outputs match the requirements, i.e., whether the Merkle root specified by the public outputs matches the Merkle tree of commitments in the ledger, whether the nullifier is not already included in the nullifier list maintained by the central bank, and whether the amount equals the transparent counter-transaction's amount. If all these requirements are satisfied, the central bank adds the new commitment and the nullifier to the ledger and notifies the client about the successful transaction.

The Merkle tree of commitments is append-only, and the mechanism that protects against double-spending, i.e., using the same old account state multiple times, is facilitated by the list of nullifiers. Consequently, the Merkle root does not necessarily have to be the most recent one. This allows users to create transactions that are accepted even if the Merkle root of the ledger has changed in the meantime due to a transaction by another user. Specifically, this option decreases the computational burden for the central bank, as it is sufficient to recompute the Merkle tree in larger epochs. Nevertheless, the epochs should not be too long, as the sequential processing of transactions by the same user requires a new Merkle tree that includes their previous commitment.

Processing this transaction, the central bank inevitably learns the amount of CBDC that is transferred from a transparent account to the privacy pool. However, since the commitments are unlinkable, it is not possible for anyone except the holder of the account to further trace these CBDC units. The same applies to the nullifier that invalidates a previous account state from the Merkle tree of commitments without revealing which specific commitment it refers to. Thus, it is impossible to determine whether a specific semi-private transaction (deposit, withdrawal) was already followed by another semi-private or fully private transaction and, particularly, whether a deposit has already been spent.

In addition to using the transparent CBDC accounts for depositing money in the privacy pool, a user could use CBDC-specific automated teller machines (ATMs) to deposit cash into the privacy pool: The user would create a transaction proposal that contains the desired amount and a ZKP similar to the depositing process described above and send it to the ATM, e.g., via Bluetooth, WiFi, or near-field communication. Then, the user inserts the corresponding amount of cash directly into the ATM. The ATM confirms the receipt of cash and forwards the user's commitment, nullifier, and ZKP to the central bank. The central bank processes the data as described above. However, the user's transparent CBDC account is not involved in this case, which provides an anonymous depositing and withdrawal process. Instead, the transaction is conducted via the CBDC account linked to the ATM, which could be provided by a PSP or the central bank directly.

**Fully private transfers**

The transfer of funds to the privacy pool enables the private bilateral exchange of CBDC. A transfer consists of two transaction proposals, i.e., individual payment instructions provided by both the sender and receiver. First, the sender and receiver need to agree on the amount of CBDC that should be transferred and on a common randomly generated number, i.e., a nonce, to increase entropy and to be able to link the two update proposals. Next, each participant creates their individual transaction proposal, including the corresponding commitment, nullifier, and ZKP, as they would in the semi-private case, with the difference that the transaction amount is not revealed but hidden by a hash (salted with the nonce). In detail, the ZKP of each of the two involved users looks as follows:

**Fully private transfer**
INPUT: Key pair, Merkle path of a previous commitment,
previous account state, (transaction) amount, nonce, role (sender/receiver)
OUTPUT: Merkle root, nullifier, new commitment, hash of amount | nonce, role
CHECKING THAT:

- *the previous commitment is contained in the tree represented by the Merkle root*
- *the previous account state belongs to the previous commitment*
- *the nullifier equals the hash of the previous account state*
- *the new account state is correct*
  *(e.g., new turnover = old turnover + amount if the role is 'sender')*
- *the new account state complies with the rules*
  *(e.g., positive balance, epoch turnover below turnover limit)*
- *the commitment equals the hash of the signed new account state*
- *the hash of amount | nonce was computed correctly*

Either the sender or the receiver batches both individual transaction proposals (including their public outputs and ZKP) and sends them to the central bank. Afterward, the central bank validates the integrity of this data through verifying the ZKP that the user generated and compares the outputs with the current state of the ledger. In particular, the central bank checks whether both Merkle roots correspond to the Merkle root of the current ledger and that the two nullifiers are not yet part of the nullifiers' list. Then, the central bank verifies whether the hashes of the value concatenated with the nonce are the same in both transactions. This check guarantees that the amount deducted from one account is equal to the amount added to the other account, without the central bank learning the actual amount due to the pre-image resistance of hash functions. Furthermore, the central bank verifies that the roles specified in the two update proposals are different, i.e., there is one sender and one receiver. Finally, the central bank adds the two new commitments to the

Merkle tree and the two nullifiers to the list and notifies the client application about the successful transfer.

Overall, the central bank only learns that a valid transaction was conducted and which two new commitments and nullifiers are involved. However, these hashes are neither related to each other in any further way, nor could they be associated with the hashes of previous or future transactions due to the unlinkability guarantees achieved through the system's construction based on commitments, nullifiers, and ZKPs. Thus, the transfers facilitated by the UAS do not provide any information about the individuals themselves, their account balances, or the amount of the transaction, and are therefore fully private.

## 3.5 Evaluation

As suggested by Hevner et al. (2004), Peffers et al. (2007), and Peffers et al. (2018), we let key stakeholders evaluate our IT artifact to assess the practical feasibility of a CBDC design that provides cash-like privacy using ZKPs. To obtain valuable insights from expert interviews, we follow the recommendations for conducting qualitative interviews by Myers and Newman (2007). In this context, we minimize social dissonances between the researcher's team and the interviewees by first introducing all interview participants to each other. To obtain a rich collection of perspectives, we talked to professionals from various backgrounds, including experts from regulation, cryptography, central banking, identity, and payments. We also made use of specific interview models, such as the waterfall technique, by first motivating our research project providing context and general definitions. Next, we provided a high-level architecture overview and a technical deep dive to improve disclosure in our interviews.

We discussed the potential risks that may occur using our CBDC with experts and demonstrated adequate mitigation measures. In addition, we presented our overall CBDC system, including the onboarding procedure, semi-private, and fully private transfers. In total, we conducted 22 interviews with a total of 44 international experts with profound expertise in the fields of law, economics, technology, and others (see Table 3.3) to reinforce the rigor of our methodological approach. Additionally, we sought to gain insights from key stakeholders regarding their key requirements for a CBDC and to receive feedback on our proposed CBDC system. Both the diversity and CBDC-specific expertise of the interviewees provided us with valuable feedback to iteratively adjust and improve our CBDC system. In addition, feedback was collected via internal discussions within the research team's organizations and presentations at various events to test the general feasibility of our approach, thereby continuously improving the artifact (Peffers et al., 2007).

During the first evaluation cycle, the experts confirmed that our design is highly innovative, that it addresses the relevant need for privacy, and that the implemented per-transaction and turnover limits on fully private transactions fit smoothly into existing regulatory frameworks from an AML and CFT compliance perspective. Many experts were surprised or even impressed by the technical capabilities of ZKPs and the maturity of the design (experts 1, 2, 3, 5, 6). In addition, the capability of our CBDC system to flexibly accommodate possible future regulatory changes, e.g., dynamically adaptable thresholds, was considered highly valuable (expert 1). Expert 2 confirmed that our approach of enforcing turnover limits on anonymous transactions would be in line with the 5th European Anti-Money Laundering Directive. The expert further stressed that tracing small-scale transactions is neither required nor desirable from a law enforcement perspective. Expert 3 emphasized that, in some jurisdictions, mistrust towards the government is considerably high, indicating that a privacy-by-design approach may be helpful to support broad adoption of a digital currency. He also pointed out that, for the usability of the payment system, it is important to reconcile and settle transparent transactions seamlessly when the limits in the privacy pool are exceeded. However, experts 4 and 5 noted that, although our adaption of the Zcash architecture is a "good trick", it may still be challenging to convince skeptical users that a solution based upon this model is, in fact, fully private. Against this background, providing the code as open-source to facilitate independent audits by cryptographers and consumer protection organizations as well as taking great educational efforts may be required to increase trust in such a cryptography-based privacy-by-design solution.

Experts 3 and 5 noted that our design could also interact smoothly with the current account-based banking systems. Moreover, expert 6 considered our approach of 'digital cash' to be intuitive, particularly illustrated through the possibility of depositing and withdrawing CBDC via an ATM. Expert 6 also mentioned that our design, in which the central bank performs the highly automatable task of verifying ZKPs and maintaining the ledger but does not need to conduct resource-intensive KYC processes, fits the roles that central banks aim for in CBDC systems (e.g., European Central Bank, 2020a).

However, the experts in the first evaluation cycle also raised two major concerns related to identity management and addressing the needs of corporations as well as those of private end-users, as they may require different limits and identity concepts. Specifically, expert 1 noted that the concept needs to be able to handle the recovery of funds in the case of lost access to the mobile phone. Also, it must provide a way to disable a privacy pool account in the case of blacklisting, e.g., if an individual is put on a sanctions list. Moreover, while our approach to provide one wallet for both a digital identity and payments was considered efficient and appealing from a user perspective (experts 6 and 7), experts 1, 6, and 7 consider the identity-based concept complex and difficult to implement in practice. According to the experts, it is unlikely that such a digital identity can be bootstrapped in the short-term, and expert 7 even considered the availability of a standardized, unique digital identity across multiple European member states a task that is as complex as introducing a CBDC itself. According to the expert, a particular challenge is that many citizens in the EU have multiple nationalities and, thus, various ID cards, which makes ID cards less suitable for guaranteeing that any citizen can only open and control one account in the privacy pool. Consequently, experts 1, 3, and 7 suggest adding intermediary-based onboarding procedures as another venue to our design.

We incorporated this feedback in our design through the following modifications:

- We added the opportunity for a joint (centralized or decentralized) ledger managed by PSPs that contains identifiers that are deterministically derived from citizens' identity, such as Hash(first name | last name | date of birth). A PSP can then sign an onboarding transaction that proves the possession of a digital certificate issued by the PSP instead of an ID card. As we already discussed in Chapter 3.4.3, the anonymity guarantees of private payments do not depend on the privacy of the onboarding process in our solution. Consequently, while the PSP learns that a specific user has registered for the privacy pool, this approach does not compromise the opportunity to conduct fully private transactions.

- We incorporated periodic proofs of non-expiration and non-revocation of the ID card (or the PSP-provided digital certificate) into our design whenever the epoch reset is performed. This modification ensures that once per epoch the user needs to prove that their ID card is still valid and hence allows to block accounts connected to an ID card that has already been revoked, e.g., in the event of loss or theft or the inclusion of an individual on a sanctions list.

In the next cycle, the experts in interview 7 pointed out that metadata, such as IP addresses, needs to be taken into account for analyzing privacy, and hence that pseudonymization is not sufficient to ensure privacy, which confirms our path of perfect unlinkability of transactions. They further noted that the *verifiability* of transactions is an important feature, i.e., a user should have the opportunity to prove that a payment did indeed happen, e.g., in the case of a lawsuit. In fact, our design already provides this capability, as a user stores their account history on their local device and can consequently reveal two consecutive previous account states in combination with the transaction confirmation signed by the central bank to demonstrate that a payment was indeed conducted with the claimed details. Furthermore, the bilateral communication between the sender and recipient preceding a transaction, where they agree upon a transaction amount and a transaction ID, can also be used to make the parties accountable bilaterally if desired by the involved parties, e.g., by revealing parts of their ID cards to the counterparty. The experts also appreciated the capability of our approach to account for embargo lists that prevent a sanctioned user from registering and further using their account through periodic checks of expiration and non-revocation of their ID card.
In interview 9, one expert noted that a balance limit might not be necessary, as a transaction cannot be larger than the turnover in a specific epoch. Nevertheless, the feasibility of balance, per-transaction, and turnover limits can account for the particular needs of various regulators. For instance, if any of these limits is not required in a particular jurisdiction, our system can be implemented easily without such a limit. The experts also pointed out that reversing a transfer must be possible

if both parties agree to do so. Also, they emphasized the importance to publish the source code in order to be able to gain broad acceptance by the public and allow for auditing by consumer protection organizations. This in turn, will, ultimately increase trust in centrally-operated payment systems and also address the concerns regarding education on privacy-by-design approaches as raised by experts 4 and 5. Considering the opportunity to prove that a payment happened, the experts in interview 9 and expert 19 added that, while this proof is indeed a desirable feature, it is also crucial to implement the possibility to delete these records to prevent measures that aim to force users to reveal them (e.g., considering coercive detention or even torture).

While confirming that a limit-based approach is suitable to make fully private payments regulatorily compliant, expert 19 argued that the current limits for cash payments are relatively low, and further reducing these limits may be difficult to justify, particularly, in view of privacy-oriented regulatory norms (see Chapter 3.1) and 'shadow economies' that may appear when limits are too low to be practical. Moreover, the expert acknowledged that the compliance by design may even help justify higher limits and that, compared to the asymmetric privacy approach by Chaum et al. (2021), our design provides true cash-like privacy. Expert 20 confirmed the suitability of ZKPs for a privacy-oriented yet regulatorily compliant form of money from a technical perspective. The expert also pointed to similar approaches based on the Ethereum blockchain that aim to create private, account-based forms of money but yet do not address regulatory constraints. He also emphasized that due to the complexity of privacy-oriented cryptographic tools, such as ZKPs, only a few stakeholders that work on CBDC are indeed aware of their technical capabilities.

An issue that was already briefly mentioned by experts 2 and 3 in the first cycle and also caused intensive discussions with experts 12, 15, 18, and 19 in the second cycle refers to the integration of companies in our CBDC system. Similar to individuals, a company should not be allowed to spend large amounts of money in the privacy pool, as it could potentially evade documentation, such as avoiding tax declarations, or supporting money laundering on a large scale. Due to the large amount and size of transactions, anonymous payments should also be restricted for businesses. We discussed two different options to incorporate companies in our CBDC design: The first approach is to allow them to open a private account via a company ID, e.g., provided by tax authorities or ingrained in trade registers. Our design is sufficiently flexible in assigning other limits to companies, which could be desired due to higher transaction sizes of businesses compared to private agents, and distinguishing between receiving and spending transactions or withdrawals of a companies' privacy pool account to their transparent account. The second approach is based on the concept of asymmetric privacy considered in Chaum et al. (2021) and Tinn and Dubach (2021), tailored to business-to-customer interactions. Indeed, by extending the semi-transparent transactions to not only include deposits and withdrawals between the accounts of one single entity, it is possible to hide the identity of the buyer and only disclose the transaction amount and the receiving company. One essential advantage of this approach may be the opportunity to be able to choose whether to keep the sender's or the receiver's identity private. For example, when the purchase of a potentially embarrassing but legal product needs to be refunded, it may be desirable to hide the identity of the end-user that previously paid anonymously.

The third design cycle further confirmed the suitability of a ZKP-based system for enabling fully private transactions and aligning with regulatory constraints through enforcing limits (experts 22, 25, and 27). Besides, expert 25 emphasized that in general, ZKPs are a "perfect tool" to align privacy and compliance requirements, but also acknowledged that, from a cryptographic perspective, it may be useful to use other forms of ZKPs, such as zk-STARKs (Ben-Sasson et al., 2018), in a potential implementation, specifically as they are regarded post-quantum secure. Expert 28 highlighted that a flexible design that allows for remuneration of CBDC is desirable from a central bank perspective, which illustrates that basic programmability features based on our account-based design are advantageous. Expert 23 also acknowledged that, in contrast to solutions such as the ECB's anonymity vouchers (European Central Bank, 2019b), our design can provide true, cash-like privacy. On the other hand, expert 22 expressed concerns that the presumably low limits for fully private transactions might imply that, despite a small share of transactions being fully private, most transactions will eventually be transparent, and hence privacy is not considerably improved overall.

One focus of this design cycle was the mitigation of risks that can potentially arise

from our design, such as ensuring that no new money can be created by private agents, that counterfeiting transactions is prevented, and that the effectiveness of the implemented limits is guaranteed. On the one hand, expert 22 warned that criminals could abuse the fully private payment system by getting access to several user accounts by purchasing accounts from other users on a black market or via blackmailing. In such a case, the effective limits could be circumvented by possessing a considerable number of accounts in the privacy pool. Although these problems can be mitigated through digital IDs that are bound to secure hardware (expert 22) and connecting the ID to other ID systems is a "great idea" (expert 23), expert 22 still pointed out that the use of secure hardware for storing keys conflicts with recovery capabilities. Experts 28 and 29 employed with a central bank also raised security concerns, particularly related to the high level of obfuscation of transaction-related information inside the privacy pool. Specifically, central banks require strong guarantees that, even if the implementation of the ZKP has a security gap, the extent to which this gap allows illicit activities or harms the monetary system remains marginal. In this regard, they also referred to an implementation error in Zcash that was detected only in 2019 and that would have allowed to 'create money out of thin air' (Fortune, 2019).
We addressed the concerns raised in the third cycle by

- conceptualizing backup capabilities with the use of secure hardware through the precautionary creation of a transaction that withdraws all funds that a user has deposited from the privacy pool to their transparent account. The user can then store this recovery backup in the cloud, potentially in encrypted form. This transaction does not provide a proof of non-revocation and non-expiration, as this would quickly be outdated, but these proofs can be provided through an in-person visit of the PSP that manages the recovery.

- mitigating risks by

  - pointing out the all-or-nothing transferability that the combination with a digital ID bound to secure hardware allows. If a user wants to sell their private account, this implies that they would need to give away their complete digital identity, meaning that they can no longer use their digital ID card, including all products associated with it, such as digital credit cards, diplomas, or health data that are bound to this ID.

– periodically closing and clearing the privacy pool so that users need to transfer their private funds to their transparent account. This process is also helpful to improve performance, as the Merkle tree and the list of commitments do not need to grow quasi infinitely.

– suggesting a hybrid approach with moderate limits for fully private transactions that are primarily meant for customer-to-customer interactions and for semi-private transactions that are primarily meant for business-to-customer interactions.

By implementing these changes, the central bank can detect some of the most critical issues imposed by potential flaws in the implementation of ZKPs, e.g., if after closing the privacy pool, the money supply would be higher than expected. Via small limits for fully private payments, the severity of the impact of selling accounts, which cannot be excluded with certainty, can be mitigated.

As our first three design cycles raised the concern that the most fragile component of our construction is its connection to a universal digital ID, we conducted another design cycle that – besides discussing the appropriateness of our risk mitigation measures – specifically included experts on digital ID systems. Expert 30 agreed that our proposed risk mitigation strategies and the combination with a digital ID might be a useful approach, although the integration with secure hardware may be considered inflexible from an end-user perspective, as they cannot access their privacy pool from multiple devices. He also suggested checking the validity of the user's ID in every transaction by default to prevent misuse. Moreover, the experts in interview 19 employed at a federal money printer in an advanced economy confirmed that the digital ID-based approach is not only more elegant but will also be possible relatively soon as many federal money printers are working on digital IDs.
Within this group, expert 35 confirmed that our proposed backup capabilities seem suitable to recover funds when losing the mobile phone. The expert also stressed that the first digital IDs that integrate with secure hardware on users' devices will be available soon and that they are potentially the only way to efficiently impede

the theft of digital IDs or their sharing or selling on a black market. The expert also found the combination of fully private transactions between end-users and semi-private transactions between end-users and companies very suitable. Moreover, expert 35 pointed out that the security of implementing ZKPs has increased substantially due to cryptographic progress in the last years, so ZKP-related security gaps discussed previously are relatively unlikely today if state-of-the-art guidelines are considered. Further, the group noted that detecting security flaws in well-audited ZKP is significantly less promising for criminals than counterfeiting paper-based money. Finally, expert 40 confirmed that our risk mitigation design may be a promising proposal to consider for central banks and that it has no obvious shortcomings. He particularly appreciated the coupling to a hardware-bound digital identity. However, as with every design proposal for a critical infrastructure, he emphasized that a thorough risk analysis would be required that is beyond the scope of this work. Nevertheless, the expert highlighted that our proposal is well presented and visualized, presenting a solid discussion basis for more in-depth analyses, and that our ZKP-based flexible design is a promising approach for the future. Experts 41 and 42 confirmed the suitability as well, emphasizing the insufficient communication between institutions that work on CBDC designs and the state of the art in cryptographic research so that the knowledge transfer via our DSR approach is highly valuable. Furthermore, experts 41 and 42 confirmed that our design incorporates privacy by design well and poses an attractive solution from the perspective of privacy-seeking users. Experts 40 and 41 also noted that, although our software-based solution cannot provide offline payments, our proposal avoids the inherent centralization of risks that comes with a hardware-based approach that may be particularly relevant in regions that do not have local companies that develop secure hardware (expert 41). For example, there is currently no provider for secure hardware on smartphones, where the manufacturing takes place in Europe. Expert 41 also pointed out that abuse cannot be excluded completely even when using hardware-bound digital identities. For instance, users could install proxies on their smartphones that transact on behalf of criminals. Yet, abuse is considerably more complex to organize than it is today with physical cash, where cash can be stolen from ATMs or cash transports, and notes can be counterfeited.

Since no considerable needs for improvements were brought forward by the experts in the fourth design cycle, and the experts interviewed in this cycle represented diverse fields, including central banks, digital identity issuers, and academics with interdisciplinary expertise from cryptography and economics, we concluded that we have reached a high level of saturation in the design and development of our artifact (Peffers et al., 2007). The critical feedback we received from key stakeholders positively influenced the design of our CBDC system, allowing us to continuously improve our artifact and thus to ultimately answer our research question.

## 3.6 General Implications

### 3.6.1 Implications for Users

Our *privacy by design* concept using ZKPs for a privacy-preserving CBDC has several implications for users. First, privacy can be strengthened, as, for fully private payments, it is impossible for third parties to identify individuals behind any transactions. This user-centric, trustless privacy empowers individuals and reduces dependencies on third parties. Privacy protection is always guaranteed without the need to trust third parties for preserving privacy. As a consequence, in our CBDC system, PSPs cannot continue pursuing business models based on analyzing and monetizing personal data. Further, trustless privacy mitigates adverse effects of data leaks, as leaked data cannot be assigned to a specific individual. Following Garratt and Van Oordt (2021), such trustless privacy can improve welfare.
Second, our proposal of combining ZKPs with digital identities supports the ongoing decentralization of business processes driven by blockchain technology and can ensure full privacy and compliance both in centralized and decentralized settings. While the necessity of using digital IDs to enable full privacy might seem contradictory to the idea of anonymity, a high level of privacy can be ensured, as no personal details of a digital ID are shared but only the *proof* that a user owns a valid ID.

### 3.6.2 Implications for Regulatory Authorities

In contrast to cash, our CBDC system enables users to reveal their complete payment history. Users are able to prove that they conducted (or did not conduct) specific transactions without revealing confidential information about individual transactions. Such proofs can serve as evidence in court because, in the case of an unjustified accusation, users can prove their innocence. Also, transaction information can be selectively disclosed according to the needs of the regulator.

Furthermore, our proposal allows for a novel, efficient, and highly automated form of *user-centric AML compliance checks*. Instead of third parties, citizens guarantee compliance *themselves* by proving possession of a valid ID and compliance with pre-defined AML rules, e.g., proving compliance with per-transaction and balance limits on anonymous payments. The PSP is not involved in the AML monitoring; it only checks the proofs provided by users, which is an easy and resource-saving task. Transaction costs can be reduced by cutting out intermediaries and eliminating resource-intensive process steps. Further, automation of compliance checks can be increased by programming and enforcing pre-defined compliance rules.

### 3.6.3 Implications for Centralized and Decentralized Use Cases

Our concept can also be applied to private-sector payment systems. While ZKPs are already used for some privacy-oriented cryptocurrencies (see Chapter 3.4.1), these privacy coins operate outside regulated domains. Our threshold approach that supports fully private and transparent payments could be applied to private-sector payment schemes to comply with AML regulation, e.g., to 'classical' cryptocurrencies and stablecoins. Such stablecoin solutions also include tokenized commercial bank money, a stablecoin fully backed by bank deposits and issued by regulated banks, and a synthetic CBDC, a stablecoin fully backed by central bank money.

Our concept can also be used for applications outside the payment domain that seek to balance privacy and compliance requirements. For example, systems for documenting and exchanging carbon certificates between organizations and/or individuals will likely require high privacy guarantees, as well.

As our concept can be used for both centralized and decentralized use cases, it can also help preserve privacy for blockchain applications. Our approach introduces the possibility for decentralized applications to comply with certain requirements without the need to share confidential data with third parties. So far, existing blockchain-based solutions for managing and exchanging tokens have not focused on stakeholders' or even regulatory privacy requirements. Regulation has rather targeted service providers that offer access to token-related services instead of end-users directly. Our approach of leveraging ZKPs makes ZKPs to some extent digital substitutes for physical hardware-based approaches without a single point of failure to achieve integrity and compliance. With a ZKP-based approach, personal data does not need to leave an end-user's device, and yet the correct local accounting is ensured by providing cryptographic proofs. A security issue in a well-audited cryptographic protocol seems less likely than a successful attack on a single secure hardware device.

### 3.6.4 Implications for Monetary Policy

Existing forms of money — cash, bank deposits, and cryptocurrencies — feature different degrees of privacy and compliance (see Figure 3.6). Cash, for example, allows for fully private transactions because information about the payer and payee — and the transaction amount — is only observable for the two transaction parties involved. To prevent illicit payments on a large scale, regulators have imposed limits on fully private cash payments so that cash can be considered compliant. Bank deposits do not allow private transactions. Transaction information is accessible for banks involved in the transaction, who can follow the money flows. Regulators imposed tight regulations for banks regarding AML and CFT, aiming at detecting illicit activities by monitoring money flows.

**Figure 3.6:** Privacy and compliance of different forms of money

| | **Full privacy** | **Compliance** |
|---|---|---|
| **Cash** | ● | ● |
| **Bank deposits** | ○ | ● |
| **Cryptocurrencies and stablecoins** | ○ | ◐ |
| **Privacy coins** | ● | ○ |
| **Our CBDC proposal** | ● | ● |

*Note*: If the dot is filled black (white), it indicates that the attribute is (not) fulfilled. If the dot is half-filled the attribute is partially fulfilled and explanation is provided in the text.

Cryptocurrencies and stablecoins do not provide a high degree of privacy, as well. Information about cryptocurrency and stablecoin transactions is accessible publicly

on open blockchain infrastructures. Even if ownership is not linked to personal information, such as name or email address, but pseudonyms, transaction information is generally visible for the general public, including regulatory agencies. Compliance, however, cannot always be guaranteed. For cryptocurrency and stablecoin transactions conducted via third parties, such as crypto exchanges, strict regulation is currently established similar to AML and CFT regulations on payments via bank deposits. However, as cryptocurrencies can also be stored locally (and offline) without publically knowing the owner, it is not possible to fully regulate such payments.
Privacy coins use privacy-preserving techniques, such as ZKPs, to conceal the identity of the transaction parties and the money flows. While providing full privacy, these cryptocurrencies are not regulated —AML and CFT requirements do not apply to privacy coins (see Chapter 3.1). For privacy coins, the trade-off between full privacy and compliance becomes visible again, as it is not possible to apply conventional AML and CFT monitoring to anonymous payment methods. The only way to restrict the use of anonymous digital means of payment is to implement limits on such payments. While per-transaction limits can easily be implemented with privacy coins, it is not possible to effectively implement turnover and balance limits, as users can open an indefinite amount of accounts (see Chapter 3.4). As a consequence, we propose to link CBDC payments to a digital ID to guarantee that every user is only allowed to open one account, allowing for the effective implementation of limits on anonymous payments.
Turning to CBDC, the ECB's key drivers to consider issuing a CBDC are, amongst others, (i) the declining use of cash as a means of payment in the euro area and (ii) the increased competition from private non-fiat denominated forms of money, such as cryptocurrencies and stablecoins (European Central Bank, 2020a). To compensate for the declining use of cash, a CBDC should mimic the properties of cash as close as possible. These properties, amongst others, include anonymous payments and regulatory compliance. Like cash and bank deposits, a CBDC has to comply with regulations around AML and CFT (European Central Bank, 2019b) because it is public money backed by the central bank and legal tender. Even if, due to the early nature of the ECB's CBDC efforts, there are no CBDC regulations in place yet, it is required that such regulations will be established prior to a CBDC introduction.

High privacy is, from our perspective, another essential feature for a CBDC — however, being more controversial due to its adverse impact on compliance measures. To date, in the context of payments, the private sector does not offer any *digital* fully private and compliant form of money (see Figure 3.6). Private sector entities do not have any incentives to give up the 'treasure' of personal data. While citizens assign high priority to privacy (European Central Bank, 2021a; European Union Agency for Fundamental Rights, 2020), in practice, they often do not behave in a privacy-preserving manner. This so-called *privacy paradox* suggests that privacy has to be promoted and secured by the *public sector* to protect citizens, e.g., by providing a privacy-preserving CBDC. If a CBDC indeed provides full privacy and regulatory compliance, it would be a close substitute for cash.

A fully private and compliant CBDC would not necessarily compete with cryptocurrencies and stablecoins. Cryptocurrencies and stablecoins are primarily used as speculative assets, for use cases around blockchain technology, e.g., in the context of Decentralized Finance (mainly Ether), or as fiat-independent stores of value (mainly Bitcoin). As cryptocurrencies exhibit high price volatility and do currently not fulfill the functions of money, the substitutability between a privacy-preserving and compliant CBDC and cryptocurrencies can be expected to be marginal. In addition, the previously described major differences in privacy and compliance point to different application domains and use cases and, thus, support the hypothesis of limited substitutability.

Users of privacy coins mainly seek a private form of money that is independent of the state and central banks, and provides the opportunity for unlimited anonymous payments. Even if a CBDC would also guarantee a high degree of privacy until a specific threshold, significant use of such a CBDC of previous privacy coin users seems unlikely.

Citizens with strong preferences for privacy could decide to switch from bank deposits to CBDC. As discussed in Chapter 2.1, a CBDC, in general, combines features of physical cash and bank deposits: It is central bank money available to private agents in digital form. However, for most CBDC proposals, the advantages of using CBDC instead of bank deposits remain unclear. In practice, except in times of crises (see Chapter 2), private agents do not consider the difference between commercial bank money and central bank money important (Deutsche Bundesbank, 2021). Even if central bank money is less risky, commercial bank money is, in many advanced economies such as Germany, secured by deposit insurance schemes and central banks act as lender of last resort, thereby reducing the underlying risk for users. High privacy guarantees, however, could drive a strong substitution of commercial bank money with CBDC as CBDC seems more attractive. Strong preferences for privacy would then imply that users could convert a substantial amount of bank deposits in CBDC. This conversion could lead to higher refinancing costs and potential liquidity shortages for banks (see Chapter 2). In our CBDC proposal, the relative attractiveness of CBDC and, thus, also the potential adverse impact on the financial sector is restricted by the limits on fully private payments. Private agents have the highest incentive to use CBDC if they can remain fully private. If the limits are exceeded and personal data is collected, CBDC's attractiveness declines. To summarize, a fully private and compliant CBDC would primarily compete with cash and bank deposits. The closer the CBDC is designed to mimic physical cash, the higher the substitution and competition with cash, and, interestingly, also with bank deposits due to privacy benefits. One key conclusion is that the size of the limits has a substantial impact on the competition of CBDC with other forms of money.[44] If limits are low, e.g., fully private CBDC payments would only be possible until 1000 euros/month, the demand for CBDC would obviously be lower than in a situation with higher limits, e.g., 10.000 euros/month.

From a monetary policy perspective, if the CBDC was a perfect substitute for cash, the composition of the liability side of the central bank's balance sheet would change. For example, a decrease in the cash position would drive a corresponding increase in the CBDC position. The total size of the central bank's balance sheet would remain constant. If the CBDC, however, was a perfect substitute for bank deposits, the balance sheet of commercial banks would, ceteris paribus, decrease, and the central bank's balance sheet would increase. This effect could imply a strong impact on

[44]Note that limits on anonymous payments can be expressed as turnover limits, balance limits, and per-transaction limits and are set by regulators.

monetary policy (Bindseil, 2020; Meaning et al., 2021).[45]

Besides the degrees of privacy and risk, there are further attributes that impact the competition between CBDC and other forms of money, and ultimately, monetary policy, which are outside the scope of this chapter. These features include, amongst others, convenience, costs, security, and remuneration. Remuneration, in particular, can be a useful tool to incentivize or disincentivize the accumulation of CBDC. As discussed in Chapter 2, a remunerated CBDC would constitute a novel monetary policy tool for central banks, which can be used to steer inflation and economic activity (Bindseil, 2020; Meaning et al., 2021; Barrdear and Kumhof, 2021; Ferrari et al., 2020).

To assess the practical relevance of full privacy and compliance for CBDC, and the impact on monetary policy, the designs of current CBDC pioneers, namely the Sand Dollar (Bahamas), the e-CNY Pilot (China), the DCash Pilot (Eastern Caribbean Currency Union), and the eNaira Pilot (Nigeria) are studied. All CBDCs feature different degrees of privacy and users can transact without sharing personal information, such as name, address, and ID information (Eastern Caribbean Central Bank, 2022; Central Bank of The Bahamas, 2021a; Central Bank of The Bahamas, 2021b; Central Bank of Nigeria, 2022; Turrin, 2022). However, a mobile number is always, and for all four CBDC projects, required. Thus, the most 'privacy-preserving' option still links all transactions to a mobile phone number. In some countries, buying a SIM card requires providing personal information so that even linking transactions to personal information cannot be ruled out. As a consequence, these payment options do not provide full privacy. In fact, all transaction data is shared with PSPs. Operators can decide to block or restrict specific accounts, which is not possible for fully private transactions. Additionally, PSPs can monitor users' transactions and report suspicious actions. Thus, compliance is implemented similar to digital payments via the financial sector.

To provide concrete examples, if choosing to use the most private, phone-number-linked payment option, users are allowed to transact, in the Bahamas, 1500 Ba-

[45] The impact of a privacy-preserving and compliant CBDC on monetary policy transmission channels is not inside the scope of this thesis. For a general analysis of the impact of CBDC on monetary policy transmission channels, see Mancini-Griffoli et al. (2018) and Meaning et al. (2021).

hamian dollars/month (Central Bank of The Bahamas, 2021a), in China, 5000 yuans/day and 50.000 yuans/year (Turrin, 2022), and in Nigeria 20,000 naira/day (Central Bank of Nigeria, 2022). In euro, this would equal in the Bahamas and in Nigeria approximately 16,000 euros/year and in China 7000 euros/year.[46]

To summarize, current CBDC pioneers do not provide fully private payments to their users. Even if officially no personal data is shared with PSPs, the mobile number is shared, and PSPs can access transaction details. Limits on these 'most private' options are relatively low, indicating that, even if users perceive these options as being fully private, the substitution between CBDC and bank deposits is low so that the CBDC mainly competes with cash. As a consequence, implications for monetary policy in these jurisdictions are relatively small.

## 3.7 Conclusion

In this paper, we followed a DSR approach to develop and evaluate an unspent-account-state-based CBDC system. To this end, we make use of recent advancements in cryptography, especially related to ZKPs. Contrary to common beliefs (e.g. Auer and Boehme, 2021; Armelius et al., 2021), we demonstrate that a software-based CBDC system can support full privacy for small transactions while addressing constraints related to AML and CFT regulation by imposing limits on anonymous payments. In our system, both privacy and compliance are provided *by design*, i.e., end-users do not have to trust third parties for preserving privacy and conducting compliance checks, as transaction data is stored only on end-users' devices (trustless privacy). We assess the feasibility and suitability of our technical artifact in 22 interviews with 44 leading experts from various fields, including regulation, computer science, cryptography, central banking, and payments.

Our artifact provides a starting point that balances privacy and compliance and may provide many avenues for future research. So far, we have presented our design to experts and instantiated the core logic in a technical implementation. Expert 22 also stated that it is important to focus on key requirements, as most design criteria tend to have trade-offs and that it is almost impossible in one research project to implement interfaces to end-users and businesses. Thus, our focus on full (cash-

[46] The calculations are based on the exchanges rates on January 27, 2022.

like) privacy and regulatory compliance is a reasonable first step. However, there are additional important CBDC design dimensions, including security, scalability, and cost, that have to be considered. Thus, there is a need for future research for a rigorous evaluation of the extent to which our IT artifact can potentially also address these other important CBDC design dimensions. Specifically, besides more detailed analyses on performance, future analyses may involve other aspects related to user experience. For example, the implications of the added complexity of a two-tiered approach, the limited recovery options, and the integration of multiple devices require innovative solutions that shield complexity from the user and survey their impact on usability. Finally, future research could study the interplay of our design with potential extensions, such as using secure hardware for facilitating offline payments, that may, however, potentially come with restricted privacy guarantees to mitigate the related security challenges.

# References


Aastveit, K., Natvik, G. J., & Sola, S. (2013). Economic uncertainty and the effectiveness of monetary policy. *Norges Bank Working Paper*, *17*.

Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, *54*(2), 442–492.

Adjemian, S., Bastani, H., Juillard, M., Karamé, F., Maih, J., Mihoubi, F., Perendia, G., Pfeifer, J., Ratto, M., & Villemot, S. (2011). Dynare Reference Manual Version 4. *Dynare Working Paper*, *1*.

Albulescu, C., Goyeau, D., & Pépin, D. (2013). Financial instability and ECB monetary policy. *Economics Bulletin*, *33*(1), 388–400.

Andolfatto, D. (2021). Assessing the impact of central bank digital currency on private banks. *The Economic Journal*, *131*(634), 525–540.

Armelius, H., Claussen, C. A., & Hull, I. (2021). On the possibility of a cash-like CBDC. *Staff memo*.

Auer, R., & Boehme, R. (2020). The technology of retail central bank digital currency. *BIS Quarterly Review, March 2020*, 85–100.

Auer, R., & Boehme, R. (2021). Central bank digital currency: The quest for minimally invasive technology. *BIS Working Papers*, *948*.

Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics*, *131*(4), 1593–1636.

Bank for International Settlements, Bank of Canada, European Central Bank, Bank of Japan, Sveriges Riksbank, Swiss National Bank, Bank of England, & Federal Reserve Board of Governors. (2020). Central bank digital currencies: Foundational principles and core features. *Series of collaborations from a group of central banks*, *1*.

Bank of Canada. (2020). Contingency planning for a central bank digital currency [https://www.bankofcanada.ca/2020/02/contingency-planning-central-bank-digital-currency/ (acess: July 11, 2021)].

Bank of England. (2020). Central bank digital currency: Opportunities, challenges and design. *Bank of England Discussion Paper*.

Barrdear, J., & Kumhof, M. (2021). The macroeconomics of central bank digital currencies. *Journal of Economic Dynamics and Control*, *104148*.

Bech, M. L., & Garratt, R. (2017). Central bank cryptocurrencies. *BIS Quarterly Review, September 2017*, 55–70.

Belke, A., & Klose, J. (2010). (How) do the ECB and the Fed react to financial market uncertainty: The Taylor rule in times of crisis. *Discussion Papers of DIW Berlin*, *972*.

Ben-Sasson, E., Bentov, I., Horesh, Y., & Riabzev, M. (2018). Scalable, transparent, and post-quantum secure computational integrity. *IACR Cryptology ePrint Archive*.

Ben-Sasson, E., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., & Virza, M. (2014). Zerocash: Decentralized anonymous payments from Bitcoin. *IEEE Symposium on Security and Privacy*, 459–474.

Ben-Sasson, E., Chiesa, A., Genkin, D., Tromer, E., & Virza, M. (2013). SNARKs for C: Verifying program executions succinctly and in zero knowledge. *Annual Cryptology Conference*, 90–108.

Bindseil, U. (2020). Tiered CBDC and the financial system. *ECB Working Paper*, *2351*.

Biryukov, A., Khovratovich, D., & Pustogarov, I. (2014). Deanonymisation of clients in Bitcoin P2P network. *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, 15–29.

Bleich, D., Fendel, R., & Rülke, J.-C. (2013). Monetary policy and stock market volatility. *Bundesbank Discussion Paper*, *45/2013*.

Boar, C., & Wehrli, A. (2021). Ready, steady, go? – Results of the third BIS survey on central bank digital currency. *BIS Papers*, *114*.

Bok, S. (1989). *Secrets: On the ethics of concealment and revelation*. Random House Digital, Inc.

Bontekoe, T. (2020). Balancing privacy and accountability in digital payment methods using zk-SNARKs [http://essay.utwente.nl/83617/1/Bontekoe_MA_EEMCS.pdf (access: June 27, 2021)].

Brainard, L. (2021). Private money and central bank money as payments go digital: An update on CBDCs [https://www.federalreserve.gov/newsevents/speech/brainard20210524a.htm (access: February 11, 2022)].

Brunnermeier, M. K., & Niepelt, D. (2019). On the equivalence of private and public money. *Journal of Monetary Economics*, *106*, 27–41.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, *12*(3), 383–398.

Camenisch, J., Hohenberger, S., & Lysyanskaya, A. (2006). Balancing accountability and privacy using e-cash. *International Conference on Security and Cryptography for Networks*, 141–155.

Castro, V. (2011). Can central banks' monetary policy be described by a linear (augmented) Taylor rule or by a nonlinear rule? *Journal of Financial Stability*, *7*(4), 228–246.

Central Bank of Nigeria. (2022). About eNaira [https://enaira.gov.ng/wallet/get-started (access: January 27, 2022)].

Central Bank of The Bahamas. (2021a). Sand dollar: Individual [https://www.sanddollar.bs/individual (access: December 17, 2021)].

Central Bank of The Bahamas. (2021b). Sand dollar: Merchants [https://www.sanddollar.bs/merchants (access: January 27, 2022)].

Chaum, D. (1983). Blind signatures for untraceable payments. *Advances in Cryptology*, 199–203.

Chaum, D., Grothoff, C., & Moser, T. (2021). How to issue a central bank digital currency. *SNB Working Papers*, *03/21*.

Cheng, J., Lawson, A. N., & Won, P. (2021). Preconditions for a general-purpose central bank digital currency [https://www.federalreserve.gov/econres/notes/feds-notes/preconditions-for-a-general-purpose-central-bank-digital-currency-20210224.htm (access: June 27, 2021)].

Chiu, J., Davoodalhosseini, S. M., Hua Jiang, J., & Zhu, Y. (2019). Central bank digital currency and banking. *Bank of Canada Staff Working Paper*, *2019-20*.

Choi, K. J., Henry, R., Lehar, A., Reardon, J., & Safavi-Naini, R. (2021). A proposal for a Canadian CBDC. *SSRN*.

Coinmarketcap. (2021). Today's cryptocurrency prices by market cap [https://coinmarketcap.com/ (access: December 14, 2021)].

Cuaresma, J. C., & Slacik, T. (2009). On the determinants of currency crises: The role of model uncertainty. *Journal of Macroeconomics*, *31*(4), 621–632.

Del Negro, M., Giannone, D., Giannoni, M. P., & Tambalotti, A. (2019). Global trends in interest rates. *Journal of International Economics*, *118*, 248–262.
Demirgüç-Kunt, A., & Detragiache, E. (2002). Does deposit insurance increase banking system stability? An empirical investigation. *Journal of Monetary Economics*, *49*(7), 1373–1406.
Deutsche Bundesbank. (2021). What do households in Germany think about the digital euro? First results from surveys and interviews. *Monthly Report*, *October 2021*, 65–84.
Diamond, D. W., & Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, *91*(3), 401–419.
Dinev, T., Xu, H., Smith, J. H., & Hart, P. (2013). Information privacy and correlates: An empirical attempt to bridge and distinguish privacy-related concepts. *European Journal of Information Systems*, *22*(3), 295–316.
Dold, F. (2019). The GNU Taler system: Practical and provably secure electronic payments. *Doctoral Dissertation, Université Rennes.*
Dyson, B., & Hodgson, G. (2016). Digital cash: Why central banks should start issuing electronic money. *Positive money.*
Eastern Caribbean Central Bank. (2022). Dcash: About the project [https://www.eccb-centralbank.org/p/about-the-project (access: January 27, 2022)].
European Central Bank. (2001). Issues related to monetary policy rules. *ECB Monthly Bulletin, October 2001*, 37–51.
European Central Bank. (2010). A real-time database of the euro area. *ECB Monthly Bulletin, September 2010*, 71–73.
European Central Bank. (2017). The future of monetary policy frameworks [https://www.ecb.europa.eu/press/key/date/2017/html/ecb.sp170525.en.html (access: October 15, 2019)].
European Central Bank. (2019a). The definition of price stability [https://www.ecb.europa.eu/mopo/strategy/pricestab/html/index.en.html (access: October 15, 2019)].
European Central Bank. (2019b). Exploring anonymity in central bank digital currencies. *IN FOCUS*, *4*.
European Central Bank. (2020a). Report on a digital euro [https://www.ecb.europa.eu/euro/html/digitaleuro-report.en.html (access: June 27, 2021)].
European Central Bank. (2020b). Study on the payment attitudes of consumers in the euro area [https://www.ecb.europa.eu/pub/pdf/other/ecb.spacereport202012~bb2038bbb6.en.pdf?05ce2c97d994fbcf1c93213ca04347dd (access February 11, 2022)].
European Central Bank. (2021a). ECB publishes the results of the public consultation on a digital euro [https://www.ecb.europa.eu/press/pr/date/2021/html/ecb.pr210414~ca3013c852.en.html, (access: June 27, 2021)].
European Central Bank. (2021b). Eurosystem launches digital euro project [https://www.ecb.europa.eu/press/pr/date/2021/html/ecb.pr210714~d99198ea23.en.html (access: July 17, 2021)].
European Convention. (2000). Charter of fundamental rights of the European Union [https://www.europarl.europa.eu/charter/pdf/text_en.pdf (access: June 27, 2021)].
European Court of Human Rights. (1950). European convention of human rights [https://www.echr.coe.int/documents/convention_eng.pdf (access: June 27, 2021)].

European Union. (2018). Directive (EU) 2018/843 of the European parliament and of the council [https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32018L0843&from=EN (access: June 27, 2021)].

European Union Agency for Fundamental Rights. (2020). Your rights matter: Data protection and privacy - fundamental rights survey [https://fra.europa.eu/en/publication/2020/fundamental-rights-survey-data-protection (access: October 5, 2021)].

Eurostat. (2020). Government finance statistics [https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Government_finance_statistics (access: October 26, 2020)].

Fauzi, P., Meiklejohn, S., Mercer, R., & Orlandi, C. (2019). Quisquis: A new design for anonymous cryptocurrencies. *International Conference on the Theory and Application of Cryptology and Information Security*, 649–678.

Feenstra, R. C. (1986). Functional equivalence between liquidity costs and the utility of money. *Journal of Monetary Economics*, *17*(2), 271–291.

Fendel, R. M., & Frenkel, M. R. (2006). Five years of single european monetary policy in practice: Is the ECB rule-based? *Contemporary Economic Policy*, *24*(1), 106–115.

Fernandez, C., Ley, E., & Steel, M. F. J. (2001a). Benchmark priors for Bayesian model averaging. *Journal of Econometrics*, *100*(2), 381–427.

Fernandez, C., Ley, E., & Steel, M. F. (2001b). Model uncertainty in cross-country growth regressions. *Journal of Applied Econometrics*, *16*(5), 563–576.

Fernández-Villaverde, J., Sanches, D., Schilling, L., & Uhlig, H. (2020a). Central bank digital currency: Central banking for all? *Review of Economic Dynamics*, *41*, 225–242.

Fernández-Villaverde, J., Sanches, D., Schilling, L., & Uhlig, H. (2020b). Central bank digital currency: When price and bank stability collide. *SSRN*.

Ferrari, M., Mehl, A., & Stracca, L. (2020). Central bank digital currency in an open economy. *CEPR Discussion Paper*, *15335*.

Ferrero, A., Gertler, M., & Svensson, L. E. O. (2008). Current account dynamics and monetary policy. *NBER Working Paper Series*, *13906*.

Fortune. (2019). Zcash discloses vulnerability that could have allowed 'infinite counterfeit' cryptocurrency [https://fortune.com/2019/02/05/zcash-vulnerability-cryptocurrency/ (access: June 27, 2021)].

Foster, D. P., & George, E. I. (1994). The risk inflation criterion for multiple regression. *The Annals of Statistics*, 1947–1975.

Friedman, M. (1963). *Inflation: Causes and consequences*. Asia Publishing House.

Garman, C., Green, M., & Miers, I. (2016). Accountable privacy for decentralized anonymous payments. *International Conference on Financial Cryptography and Data Security*, 81–98.

Garratt, R. J., & Van Oordt, M. R. (2021). Privacy as a public good: A case for electronic cash. *Journal of Political Economy*, *129*(7).

Gerdesmeier, D., & Roffia, B. (2003). Empirical estimates of reaction functions for the euro area. *ECB Working Paper Series*, *206*.

Gerlach, S. (2007). Interest rate setting by the ECB, 1999-2006: Words and deeds. *International Journal of Central Banking*, *3*(3), 1–46.

Gertler, M., & Karadi, P. (2011). A model of unconventional monetary policy. *Journal of Monetary Economics*, *58*(1), 17–34.

Goldwasser, S., Micali, S., & Rackoff, C. (1989). The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, *18*(1), 186–208.

Gregor, S., & Hevner, A. (2013). Positioning and presenting design science research for maximum impact. *MIS Quarterly*, *37*, 337–355.

Gross, J., & Schiller, J. (2021). A model for central bank digital currencies: Implications for bank funding and monetary policy. *SSRN Electronic Journal*.

Gross, J., Sedlmeir, J., Babel, M., Bechtel, A., & Schellinger, B. (2021). Designing a central bank digital currency with support for cash-like privacy. *SSRN*.

Gross, J., & Zahner, J. (2021). What is on the ECB's mind? Monetary policy before and after the global financial crisis. *Journal of Macroeconomics*, *68*, 103292.

Grothoff, C., & Dold, F. (2021). Why a digital euro should be online-first and bearer-based [https://taler.net/papers/euro-bearer-online-2021.pdf (access: July 18, 2021)].

Guerrieri, L., & Iacoviello, M. (2015). Occbin: A toolkit for solving dynamic models with occasionally binding constraints easily. *Journal of Monetary Economics*, *70*, 22–38.

Hayo, B. (2018). On standard-error-decreasing complementarity: Why collinearity is not the whole story. *Journal of Quantitative Economics*, *16*(1), 289–307.

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, (1), 75–105.

Hollo, D., Kremer, M., & Duca, M. (2012). CISS – A composite indicator of systemic stress in the financial system. *ECB Working Paper Series*, *1426*.

Iivari, J. (2015). Distinguishing and contrasting two strategies for design science research. *European Journal of Information Systems*, *24*(1), 107–115.

Kaefer, B. (2014). The Taylor rule and financial stability – A literature review with application for the eurozone. *Review of Economics*, *65*(2).

Kahn, C. M. (2018). Payment systems and privacy. *SSRN*.

Kahn, C. M., McAndrews, J., & Roberds, W. (2005). Money is privacy. *International Economic Review*, *46*(2), 377–399.

Kannengießer, N., Lins, S., Dehling, T., & Sunyaev, A. (2020). Trade-offs between distributed ledger technology characteristics. *ACM Computing Surveys*, *53*(2).

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, *90*(430), 773–795.

Keister, T., & Sanches, D. (2019). Should central banks issue digital currency? *FRB Philadelphia Working Paper*, *19-26*.

Kiff, J., Alwazir, J., Davidovic, S., Farias, A., Khan, A., Khiaonarong, T., Malaika, M., Monroe, H., Sugimoto, N., Tourpe, H., & Zhou, P. (2020). A survey of research on retail central bank digital currency. *IMF Working Paper*, *20/104*.

Koop, G. M. (2003). *Bayesian econometrics*. John Wiley & Sons Inc.

Kubach, M., & Sellung, R. (2021). On the market for self-sovereign identity: Structure and stakeholders. *Open Identity Summit 2021*.

Lane, P. (2019). Wir wollen positive Zinsen. *Handelsblatt: Düsseldorf*, *187*, 26–27.

Lane, T. (2020). Money and payments in the digital age [https://www.bis.org/review/r200311d.pdf (access: June 27, 2021)].

Lee, K., Morley, J., & Shields, K. (2015). The meta Taylor rule. *Journal of Money, Credit and Banking*, *47*(1), 73–98.

Lee, K., Olekalns, N., & Shields, K. (2013). Meta Taylor rules for the UK and Australia; accommodating regime uncertainty in monetary policy analysis using model averaging methods. *The Manchester School*, *81*, 28–53.

Ley, E., & Steel, M. F. J. (2009). On the effect of prior assumptions in Bayesian model averaging with applications to growth regression. *Journal of Applied Econometrics*, *24*(4), 651–674.

Mancini-Griffoli, T., Peria, M. S. M., Agur, I., Ari, A., Kiff, J., Popescu, A., & Rochon, C. (2018). Casting light on central bank digital currency. *IMF Staff Discussion Note*, *8*.

March, S. T., & Smith. (1995). Design and natural science research on information technology. *Decision Support Systems*, *15*(4), 251–266.

Mastercard. (2021). Mastercard and Island Pay launch world's first central bank digital currency-linked card [https://www.mastercard.com/news/press/2021/february/mastercard-and-island-pay-launch-world-s-first-central-bank-digital-currency-linked-card/ (access: July 18, 2021)].

Mattke, J., Hund, A., Maier, C., & Weitzel, T. (2019). How an enterprise blockchain application in the US pharmaceuticals supply chain is saving lives. *MIS Quarterly Executive*, *18*(4).

Meaning, J., Dyson, B., Barker, J., & Clayton, E. (2021). Broadening narrow money: Monetary policy with a central bank digital currency. *International Journal of Central Banking*, *68*.

Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M., & Savage, S. (2016). A fistful of Bitcoins: Characterizing payments among men with no names. *Communications of the ACM*, *59*(4), 86–93.

Merkle, R. C. (1987). A digital signature based on a conventional encryption function. *Conference on the Theory and Application of Cryptographic Techniques*, 369–378.

Miers, I., Garman, C., Green, M., & Rubin, A. D. (2013). Zerocoin: Anonymous distributed e-cash from Bitcoin. *IEEE Symposium on Security and Privacy*, 397–411.

Milani, F. (2004). Monetary policy with a wider information set: A Bayesian model averaging approach. *SSRN*.

Milani, F. (2008). Monetary policy with a wider information set: A Bayesian model averaging approach. *Scottish Journal of Political Economy*, *55*(1), 1–30.

Molodtsova, T., & Papell, D. (2012). Taylor rule exchange rate forecasting during the financial crisis. *NBER Working Paper Series*, *18330*.

Moral-Benito, E. (2010). Model averaging in economics. *CEMFI Working Paper*, *1008*.

Myers, M., & Newman, M. (2007). The qualitative interview in IS research: Examining the craft. *Information and Organization*, *17*, 2–26.

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system [http://bitcoin.org/bitcoin.pdf (access: June 27, 2021)].

Niepelt, D. (2020). Monetary policy with reserves and CBDC: Optimality, equivalence, and politics. *EPR Discussion Paper*, *15457*.

Obstfeld, M., & Rogoff, K. S. (2005). Global current account imbalances and exchange rate adjustments. *Brookings Papers on Economic Activity*, *2005*(1), 67–146.

Odlyzko, A. (2004). Privacy, economics, and price discrimination on the internet. *Economics of Information Security* (pp. 187–211). Springer.

Orphanides, A. (2001). Monetary policy rules based on real-time data. *American Economic Review*, *91*(4), 964–985.

Panetta, F. (2018). 21st century cash: Central banking, technological innovation and digital currency. In E. Gnan & D. Masciandaro (Eds.), *Do we need central bank digital currency?* (pp. 23–33).

Panetta, F. (2021). A digital euro to meet the expectations of Europeans [https://www.ecb.europa.eu/press/key/date/2021/html/ecb.sp210414_1~e76b855b5c.en.html (access: June 27, 2021)].

Partala, J., Nguyen, T. H., & Pirttikangas, S. (2020). Non-interactive zero-knowledge for blockchain: A survey. *IEEE Access*, *8*, 227945–227961.

Peek, J., Rosengren, E. S., & Tootell, G. (2016). Does Fed policy reveal a ternary mandate? *Federal Reserve Bank of Boston Working Papers*, *16-11*.

Peffers, K., Tuunanen, T., & Niehaves, B. (2018). Design science research genres: Introduction to the special issue on exemplars and criteria for applicable design science research. *European Journal of Information Systems*, *27*(2), 129–139.

Peffers, K., Tuunanen, T., Rothenberger, M., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, *24*(3), 45–77.

Pocher, N., & Veneris, A. (2021). Privacy and transparency in CBDCs: A regulation-by-design AML/CFT scheme [https://www.eecg.utoronto.ca/~veneris/21icbc2.pdf (access: June 27, 2021)].

Roskelley, K. D. (2016). Augmenting the Taylor rule: Monetary policy and the bond market. *Economics Letters*, *144*, 64–67.

Rossi, M., Mueller-Bloch, C., Thatcher, J. B., & Beck, R. (2019). Blockchain research in information systems: Current trends and an inclusive future research agenda. *Journal of the Association for Information Systems*, *20*(9), 1388–1403.

Rotemberg, J. J. (1982). Sticky prices in the united states. *Journal of Political Economy*, *90*(6), 1187–1211.

Sander, T., & Ta-Shma, A. (1999). Flow control: A new approach for anonymity control in electronic cash systems. *International Conference on Financial Cryptography*, 46–61.

Sauer, S., & Sturm, J.-E. (2007). Using Taylor rules to understand European Central Bank monetary policy. *German Economic Review*, *8*(3), 375–398.

Sidrauski, M. (1967). Rational choice and patterns of growth in a monetary economy. *The American Economic Review*, *57*(2), 534–544.

Silfversten, E., Favaro, M., Slapakova, L., Ishikawa, S., Liu, J., & Salas, A. (2020). *Exploring the use of Zcash cryptocurrency for illicit or criminal purposes*. RAND.

Stark, J. (2010). Geld- und Fiskalpolitik während und nach der Krise [https://www.ecb.europa.eu/press/key/date/2010/html/sp101015.de.html (access: October 15, 2019)].

Sturm, J.-E., & Wollmershäuser, T. (2008). The stress of having a single monetary policy in europe. *CESifo Working Paper*, *2251*.

Svensson, L. E. (2003). What is wrong with Taylor rules? Using judgment in monetary policy through targeting rules. *Journal of Economic Literature*, *41*(2), 426–477.

Taylor, J. B. (1993). Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy*, *39*, 195–214.

Tefft, S. K. (1980). *Secrecy a cross-cultural perspective*. Human Sciences Press.

The Guardian. (2021). New Zealand's central bank says its systems have been hacked [https://www.theguardian.com/world/2021/jan/11/new-zealands-central-bank-says-its-systems-have-been-hacked (access: June 27, 2021)].

Tinn, K., & Dubach, C. (2021). Central bank digital currency with asymmetric privacy [https://www.mcgill.ca/engineering/files/engineering/central_bank_digital_

currency_with_asymmetric_privacy_mcgill_tinn_dubach.pdf (access: June 27, 2021)].
Turrin, R. (2022). Digital yuan e-CNY App review [https://watch.mmhmm.app/z_H3sDX3cvbeXjoOSAD18o (access: January 27, 2022)].
United Nations. (1948). Universal declaration of human rights [https://www.un.org/en/udhrbook/pdf/udhr_booklet_en_web.pdf (access: June 27, 2021)].
Veneris, A., Park, A., Long, F., & Puri, P. (2021). Central bank digital loonie: Canadian cash for a new global economy [https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3770024 (access: June 27, 2021)].
Wu, J. C., & Xia, F. D. (2016). Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Journal of Money, Credit and Banking*, *48*(2-3), 253–291.
Wüst, K., Kostiainen, K., & Capkun, S. (2021). Platypus: A central bank digital currency with unlinkable transactions and privacy preserving regulation.
Zcash Team. (N.d.). What are zk-SNARKs? [https://z.cash/technology/zksnarks/ (access: February 02, 2022)].
Zellner, A. (1986). On assessing prior distributions and Bayesian regression analysis with g-prior distributions. In P. Goel & A. Zellner (Eds.), *Bayesian inference and decision techniques: Essays in honor of Bruno de Finetti* (pp. 233–243). Elsevier Science Publishers.
Zeugner, S., & Feldkircher, M. (2015). Bayesian model averaging employing fixed and flexible priors: The BMS package for R. *Journal of Statistical Software*, *68*(4), 1–37.
Zhang, R., Xue, R., & Liu, L. (2019). Security and privacy on blockchain. *ACM Computing Surveys*, *52*(3), 1–34.
Zwick, D., & Dholakia, N. (2004). Whose identity is it anyway? Consumer representation in the age of database marketing. *Journal of Macromarketing*, *24*(1), 31–43.