THREE ESSAYS ON CENTRAL BANKING

A Dissertation

by

RYAN AUSTIN RHOLES

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Chair of Committee,	Catherine Eckel
Committee Members,	Luba Petersen
	Sarah Zubairy
	Marco Palma
Head of Department,	Steve Puller

May 2021

Major Subject: Economics

ABSTRACT

This dissertation comprises three chapters focused on the conduct of monetary policy. Specifically, I use a blend of empirical and experimental research methods study unconventional monetary policy alternatives, to quantify the effect of information flow on markets, and to examine the effi-cacy of emergent changes in central bank communication.

Chapter one of this dissertation introduces an experimental framework that serves as a test-bed for both monetary and fiscal policies. This framework allows policymakers and academics to generate data regarding the efficacy of policy alternatives in a relatively low-cost, safe, and controlled environment. To validate this experimental framework, we explore two policy alternatives in a constrained policy environment: allowing a central bank to use negative nominal interest rates and permanently increasing the inflation target. Our findings indicate that negative interest rates can effectively stimulate demand, consistently. closing an output and employment gaps. However, permanently increasing an inflation target cannot.

Chapters two and three of this dissertation focus on central bank communication. Chapter two uses a high-frequency identification approach to quantify the effect of information flow on financial markets in the United Kingdom. Specifically, we study whether and how the U.K.'s yield curve responds to communication from the Bank of England (BoE) about the forecast uncertainty and balance of risks surrounding both inflation and output growth. We find that yields respond at least as strongly to this higher-order information as they do to information about the BoE's expected paths for inflation and output growth. Chapter three describes an experiment in which we study how uncertainty in central bank communication influences individual expectations formation and individual forecast uncertainty. We find that reducing the precision of central bank communication by including uncertainty can reduce the coordinating power of projections and increase individual-level uncertainty.

DEDICATION

To my wife, Irina, for her love, patience, and support.

To my daughter, Sofia, who is my world. I love you forever.

To my mother, Jon. I turned out okay!

To Sugar Bear, Weasel, and Fabri. I will live a good life to honor you. Semper Fidelis.

ACKNOWLEDGMENTS

I am forever indebted to my advisor, Catherine Eckel. Catherine took me on as a young(ish), naive graduate student, helped me find my passion, and then provided me endless support in its pursuit. Working with and learning from such a truly wonderful person is an experience I will always cherish. Thank you, Catherine.

I want to acknowledge my co-author and committee member, Luba Petersen. Luba took me into the trenches of experimental research and sacrificed incredible amounts of time and sanity there to teach me. She demanded excellence and I am better for it. Thank you, Luba, for everything.

I am also grateful to Tatevik Sekhposyan and Sarah Zubairy for teaching me so much about macroeconomics and their willingness to support me in my effort to bring experiments to the field. And to Marco Palma, for his support and insight.

CONTRIBUTORS AND FUNDING SOURCES

Contributors

This work was supported by a dissertation committee consisting of Professor Catherine Eckel, the committee chair, and committee members Dr. Sarah Zuibary, Dr. Marco Palma, and Dr. Luba Petersen. The work herein was co-authored with Dr. Luba Petersen (chapters 2 and 4) and with Dr. Tatevik Sekhposyan (chapter 3).

Funding Sources

The research in this dissertation was funded by the National Science Foundation through a Doctoral Dissertation Research Improvement Grant and by an SSHRC Development and Insight Grant.

NOMENCLATURE

BoE	Bank of England

TABLE OF CONTENTS

Page

ABSTRACT ii

DEDICATION iii

ACKNOWLEDGMENTS iv

CONTRIBUTORS AND FUNDING SOURCES v

NOMENCLATURE vi

TABLE OF CONTENTS vii

LIST OF FIGURES ix

LIST OF TABLES xi

1. INTRODUCTION 1

2. ESCAPING SECULAR STAGNATION WITH UNCONVENTIONAL MONETARY POLICY 3

2.1 Introduction 3
2.2 Literature 5
2.2.1 Theoretical and Empirical 5
2.2.2 Experimental Macroeconomics 7
2.3 Theoretical Framework and Hypotheses 11
2.3.1 Households 11
2.3.2 Firms 12
2.3.3 Central Bank 13
2.3.4 Equilibrium 14
2.3.5 Monetary Policy 15
2.4 Experimental Implementation 19
2.4.1 Experiment 20
2.4.2 Treatments 28
2.4.2.1 *Baseline* 29
2.4.2.2 Policy treatments 29
2.5 Results 30
2.6 Conclusion 45
2.7 References 46

3. CENTRAL BANK DENSITY FORECASTS: DO HIGHER-ORDER MOMENTS MATTER? ... 54

3.1 Data ... 57
3.1.1 Timing ... 60
3.1.2 Density forecast data ... 60
3.1.3 Daily data for the UK's yield curve ... 64
3.1.4 Data Transformations ... 64
3.2 Results ... 67
3.2.1 Expectations ... 71
3.2.2 Historical Decomposition ... 74
3.3 Conclusion ... 76
3.4 References ... 77

4. SHOULD CENTRAL BANKS COMMUNICATE UNCERTAINTY IN THEIR PROJECTIONS?* ... 80

4.1 Introduction ... 80
4.2 Experimental Design ... 83
4.2.1 Data-generating process ... 86
4.2.2 Payoffs ... 88
4.2.3 Treatments ... 89
4.2.4 Procedures ... 92
4.2.5 Hypotheses ... 92
4.3 Results ... 99
4.3.1 Aggregate results ... 99
4.3.2 Individual results ... 102
4.4 Conclusion ... 118
4.5 References ... 121

5. SUMMARY AND CONCLUSIONS ... 126

APPENDIX A. ... 131

A.1 Instructions ... 131
A.1.1 Group Instructions ... 131
A.1.2 Individual Instructions ... 140
A.2 Comparison of two-period ahead and one-period ahead data-generating processes ... 147
A.3 Time series ... 149

REFERENCES ... 131

LIST OF FIGURES

FIGURE Page

2.1 A deleveraging shock moves an economy from a unique inflationary to a unique secular stagnation equilibrium. 16

2.2 Increasing the inflation target generates two inflationary equilibria: a target, full-employment equilibrium and a liquidity trap equilibrium. 16

2.3 Removing the ZLB removes the kink from the AD curve, thereby eliminating the secular stagnation equilibrium and creating instead a unique inflationary equilibrium. 18

2.4 Payoff function for an individual price forecast 21

2.5 Co-determination of middle-aged and old spending 23

2.6 Stage 1 Screen 24

2.7 Stage 2 Screen 25

2.8 Aggregate inflation by session for *Baseline* 32

2.9 Nominal interest rates by session for *Baseline* 33

2.10 Nominal interest rates by session for *Baseline* 34

2.11 Aggregate inflation dynamics by session for *HigherTarget* 35

2.12 Nominal interest rates by session for *HigherTarget* 37

2.13 CDFs of consumption for all subjects in all periods of *HigherTarget* 38

2.14 Aggregate inflation dynamics by session for *NegativeIR* 39

2.15 Nominal interest rates by session for *NegativeIR* 40

2.16 CDFs of consumption for all subjects and all periods for *NegativeIR* 41

2.17 Forecasting heuristics by phase and by treatment 44

3.1 Density forecasts from the Bank of England's Monetary Policy Report. 58

3.2 Policy and Communication Schedule 61

3.3 First moments of the BoE's density forecasts of GDP growth and inflation 62

3.4 Uncertainty and skewness of GDP growth and inflation 63

3.5 Sequential forecast example 65

3.6 GDP growth and inflation revision factors 66

3.7 Yield curve factor loadings 67

3.8 Yield responses to information shocks 69

3.9 Placebo estimates 72

3.10 Response of professional forecasters to information shocks 73

3.11 Historical decomposition of yield curve factor revisions 75

4.1 Flow of information, decisions, and outcomes 85

4.2 Simulated statistics under alternative forecasting heuristics 94

4.3 Time series of Group treatments 100

4.4 Time series of Individual treatments 101

4.5 Distributions of absolute forecast errors 103

4.6 Distribution of forecasting heuristics for $t+1$ inflation, by treatment 117

A.1 Aggregate statistics in two- and one-period ahead NK DGPs 147

A.2 Time series of NoComm sessions 149

A.3 Time series of Point sessions 150

A.4 Time series of Point&Density sessions 150

LIST OF TABLES

TABLE Page

2.1 Parameterized equilibria across treatments and phases 30

2.2 Inflation forecasting heuristics..... 42

3.1 Table of correlations for information shocks 65

4.1 Treatments summary 90

4.2 Summary statistics of aggregate and individual forecast variables 104

4.3 Absolute forecast errors..... 106

4.4 Disagreement in inflation forecasts..... 109

4.5 Uncertainty in inflation forecasts 111

4.6 Credibility of central bank projections 114

4.7 Forecasting heuristics 116

1. INTRODUCTION

This dissertation focuses on the conduct of monetary policy. Specifically, I use a blend of empirical and experimental research methods study unconventional monetary policy alternatives,to quantify the effect of information flow on markets, and to examine the efficacy of emergent changes in central bank communication.

The first chapter of my dissertation introduces an experimental framework that serves as a testbed for monetary policy, fiscal policy, and central bank communication. This framework is built around a simple overlapping generations model that can cleanly incorporate aggregate demand shocks and that features a rules-based mechanistic central bank. The experimental economies in this framework evolve according to the price expectations and real decisions of student subjects. Importantly, aggregate demand shocks in this framework consistently generate deflationary and liquidity trap episodes, which creates an idea environment to test the efficacy of policy interventions. We validate this experimental environment by testing two unconventional monetary policy alternatives: allowing a central bank to use negative policy rates and allowing a central bank to permanently increase its inflation target. We find that implementing negative rates can consistently pull economies out of liquidity traps, closing output and employment gaps. However, contrary to theoretical predictions, permanently increasing the central bank's inflation target cannot.

The second chapter of my dissertation considers the Bank of England's density forecasts and its revisions to quantify the effects of information flow on the financial markets and survey forecasters. Central banks have increasingly relied on published forecasts to communicate their economic outlook to market participants. Point forecasts and their revisions have been shown to move financial markets. The effects of the higher-order moments, however, have not been investigated thoroughly, and this is primarily due to data limitations. The Bank of England, on the other hand, has been publishing information on its density forecasts since the late 1990s, making it useful for our analysis. Using daily information on the financial markets, we find that the updates of higher moments are more important in moving financial markets than the revisions to the first central moment of the density forecasts, making them relevant for monetary policy communication. Information about output matters more than information about inflation, and the effect of information is state-contingent. Finally, we see that the consensus forecast and level of forecast disagreement, among professional forecasters are strongly correlated with updates in higher-order forecast moments.

This third chapter of my dissertation takes an experimental approach to provide further evidence on the emerging practice by central banks of communicating forecast uncertainty to market participants. To do this, we compare the effects of point and density projections in a learning-to-forecast laboratory experiment where participants' aggregated expectations about one- and two-period-ahead inflation influence macroeconomic dynamics. Precise point projections are more effective at managing inflation expectations. Point projections reduce disagreement and uncertainty while nudging participants to forecast rationally. Supplementing the point projection with a density forecast mutes many of these benefits. Relative to a point projection, density forecasts lead to larger forecast errors, greater uncertainty about own forecasts, and less credibility in the central bank's projections. We also explore expectation formation in individual-choice environments to understand the motives for responding to projections. Credibility in the projections is significantly lower when strategic considerations are absent, suggesting that projections are primarily effective as a coordination device. Overall, our results suggest that communicating uncertainty through density projections reduces the efficacy of inflation point projections.

2. ESCAPING SECULAR STAGNATION WITH UNCONVENTIONAL MONETARY POLICY

2.1 Introduction

Many developed economies exhibit tell-tale symptoms of secular stagnation: decades-long downward trends in natural interest rates, tepid output growth well below estimates of potential, growing debt-to-GDP ratios, negative real interest rates, and below-target inflation. Furthermore, these conditions persist despite expansionary policies employed by the Federal Reserve and by other central banks. For instance, the United States spent seven years at its zero lower bound before lifting off in 2015 and slowly climbed up to a meager 2.25% by 2019 thanks to, among other things, extraordinary quantitative easing and forward guidance. Summers (2019) argued that this was insufficient space to operate in a future recession. At the 2019 Jackson Hole Economic Policy Symposium, Federal Reserve Chairman Jerome Powell described this proximity to the effective lower bound as "the pre-eminent monetary policy challenge of our time" as central bankers sought out fresh policy tools.

Now, the unprecedented economic crisis brought on by the global pandemic has highlighted the pressing need to expand the tool kits of policy-makers. Monetary interventions once considered too extreme or perhaps outlandish are now finding their way into mainstream policy discussions. Many central bankers are now giving serious consideration to raising inflation targets and implementing negative policy rates, both of which would, in theory, stimulate inflation expectations and propel economic activity. However, policymakers are hesitant to adopt these tools given the limited evidence on their efficacy.

In this paper, we build a flexible and novel experimental environment to test-bed these unconventional policies, thereby providing policy makers with much-needed, timely evidence on their effectiveness. Importantly, this environment, though tractable, captures many realistic and pressing concerns by allowing for secular stagnation and liquidity traps at the zero lower bound. We use this framework to provide the first experimental evidence on the potency of raising the inflation target and negative nominal interest rates in combating secular stagnation.

To this end, we adapt Eggertsson, Mehrotra, and Robbins (2019, EMR henceforth) to the lab. In each experimental economy, groups of participants form expectations and make financial decisions together with automated firms and policy makers. We expose each economy to an exogenous deleveraging shock that should generate pessimistic inflation expectations and a secular stagnation of the economy.

In our *Baseline* treatment, We find that most economies converge toward the initial full-employment equilibrium and that deleveraging shocks consistently generate pessimistic expectations and typically manufacture various degrees of instability manifesting itself either in permanent deflation. In most sessions the economies converge in the direction of the secular stagnation equilibrium.

We then explore alternative policy options to return the economy to its full employment equilibrium after an extended episode at the zero lower bound. In our *HigherTarget* treatment, we raise the inflation target from 10-30%, one policy prescription formalized by EMR. We find that significantly increasing the central bank's target, which in theory yields a multiplicity of equilibria, does not effectively return any of the economies to the targeted, full-employment equilibrium. On the other hand, removing the ZLB in our *NegativeIR* treatment consistently stimulates spending, which in turn encourages inflationary expectations and a quick escape from the zero lower bound. These economies stabilize at a full-employment equilibrium that coincides with the central bank's inflation target.

We see in all treatments that deleveraging shocks generate considerable consumption heterogeneity. Thus, the decrease in welfare is larger for some subjects than others following these shocks. We also observe that deleveraging shocks effect expectations formation. Instability brings about a high level of constant gain learning and other backward-looking heuristics.

The main takeaway from our experiments is that negative interest rates are more effective than raising inflation targets at shifting consumption to the present and stimulating an economy out of secular stagnation. Our findings suggest that inflation expectations and demand is better stimulated through realized wealth effects than coordination on rational expectations equilibria. Our participants fail to perceive a central bank's higher inflation target credible after an extended period of deflation.

We also make a number of valuable methodological contributions in this paper. First, we demonstrate how a complex general equilibrium theoretical framework can be distilled to a simple implementation that nonetheless allow for a meaningful interaction of expectations, decisions, and monetary policy. This framework can flexibly be extended to allow for fiscal policy, credit markets, policy communication and coordination. Additionally, we contribute to the experimental literature by bridging learning-to-forecast and production economy experiments. Fusing these frameworks allows us to link expectations and real decision in response to policy, filling a crucial but neglected empirical gap.

The rest of the paper is organized as follows. Section 2 places our paper in the context of the existing macroeconomics and experimental research. Section 3 lays out the theoretical framework and hypotheses for our experiment. Section 4 provides details of the experimental implementation, Section 5 presents our experimental results, and Section 6 concludes.

2.2 Literature

This paper makes important contributions to the macroeconomic literature on economic growth and policy, as well as the experimental macroeconomic literature. We discuss these contributions below.

2.2.1 Theoretical and Empirical

The existing liquidity trap literature, though large, does not provide a thorough treatment of secular stagnation. As highlighted by EMR, this is because most models employ representative agent frameworks[1] where the long-run interest rate is uniquely determined by a representative discount factor. Since the real interest rate must eventually revert to a positive long-run level, ZLB episodes triggered by transitory shocks are themselves temporary. This theoretical outcome

[1] See Krugman (1998), Eggertsson and Woodford (2003), Christiano, Eichenbaum and Rebelo (2011), Eggertson and Krugman (2012), and Werning (2012) for examples.

yields a 'wait-and-see' policy approach to liquidity traps, which would constitute an obvious failure of policy for an economy facing secular stagnation. To circumvent this issue and allow for permanently-negative neutral real interest rates, EMR introduce ZLB episodes into an OLG model (This was first shown by Samuelson (1958)).

Deflationary steady states can emerge in the New Keynesian framework. Benhabib, Schmitt-Grohe and Uribe (2001), Schmitt-Grohe and Uribe (2017), and Benigno and Fornaro (2015) all feature a ZLB that binds due to hysteresis. However, ZLB steady states in these models are locally indeterminate and are open to the criticism that ZLB episodes driven by hysteresis are not 'learnable' and it is therefore unclear how such a steady state can coordinate expectations (Christiano, Eichenbaum and Johannsen, 2016). However, Arifovic, Schmitt-Grohe, and Uribe (2018) show that the liquidity trap equilibrium is learnable under social learning. Additionally, Gibbs (2018) shows mathematically that both the deflationary and targeted-inflation equilibrium in EMR's model are both E-stable when they exist and that the full-employment, liquidity trap equilibrium is only stable whenever agents in this OLG model face sufficiently large borrowing limits.

There is very limited evidence on the ability of central banks to increase inflation expectations by *raising* their inflation target. Most inflation targeting central banks have kept their inflation targets constant or reduce them after stabilizing their price level growth. There are two exceptions. The Reserve Bank of New Zealand (RBNZ) first began targeting inflation in the range of 0 to 2 percent, and were quite effective in bringing inflation down from well over 6 percent in 1990 to an average of 2.8 percent in the five years that followed. In 1996, the RBNZ increased the upper bound of their range to 3 percent, effectively raising the midpoint from 1 to 1.5 percent. From 1996 to 2002, inflation averaged 1.8 percent, converging right toward the mid-point of the target range. In 2003, the Bank raised the lower bound of the range from 0 to 1 percent, increasing the midpoint to 2 percent. Over the next five years, average inflation rose to 2.8 percent, and from 2003-2019, averaged 2 percent. More recently, the Bank of Japan (BoJ) experimented with communicating an explicit inflation target and raising its target. In February 2012, the BoJ announced that it would explicitly target inflation at 1 percent (this had been the implicit mid-point of an acceptable range of inflation since 2006). This announcement led to a modest reduction in deflation. In January 2013, BoJ further increased its target from 1 to 2 percent. While there was some rapid inflation growth over the following year, inflation has fluctuated between 0.5-1%. Overall, it appears that raising the target had a positive effect on inflation, albeit smaller than intended.

Finally, this paper contributes to an emerging literature that studies the use of negative policy rates. This literature, while very important, is woefully incomplete due to an extreme lack of data. Eggertsson, Juelsrud, Summers, and Wold (2019) use empirical data to show that negative policy rates produce a lower bound on household deposits (DLB). They embed this DLB into a banking sector model and show that negative rates are expansionary only under some conditions. Altavilla, Burlon, Giannetti, and Holton (2019) use Euro-area data to show empirically that the transmission mechanism of monetary policy does not break down when rates become negative. Wu and Xia (2016) build a shadow rate term structure model (SRTSM) calibrated to the Euro Area and show that all four of the ECB's negative rate cuts lowered maturities along the short-end of the yield curve. Further, they show that forward guidance may facilitate the transmission of negative rates.

Our experiment contributes to this literature by building a flexible framework capable of producing empirical data surrounding the use of higher inflation targets and negative policy rates. Our results mostly align with this existing literature in the sense that negative rates in our experimental economies stimulate aggregate demand yielding an expansionary effect. Importantly, our framework provides insight to how individual decisions and expectations respond to negative policy rates, thereby shedding light on a transmission channel of negative rates that real-world data alone cannot elucidate.

2.2.2 Experimental Macroeconomics

Begun largely in response to an invitation from Robert Lucas (1986), experimental macroeconomics is a relatively young but fruitful field of study focused on testing the microfoundations of modern macroeconomic models in a controlled, laboratory setting (see John Duffy (2016) for a thorough survey of the literature). Several branches of literature from this emergent field relate closely to our research.

Dynamic Optimization & Learning to Optimize in Individual and Production Economies Settings

Laboratory experiments testing the ability of people to solve dynamic optimization problems in the form of a one-sector, infinite-horizon model reveal that experimental subjects (relative to theoretical predictions) do a poor job of consumption smoothing (consumption often too closely tracks current income) and that agents' consumption decisions are not time-independent (For examples, see Hey and Dardanoni (1988), Carbone and Hey (2004), Noussair and Matheny (2000), Lei and Noussair (2002), Ballinger et al. (2003), Carbone (2006), Crockett and Duffy (2013), Carbone and Duffy (2014) and Meissner (2016)). Possible explanations for such failures are binding liquidity constraints, precautionary saving, and debt aversion. There are some exceptions, however. Crockett and Duffy (2013) finds that subjects facing a concave utility function are able to optimize within the framework of a Lucas Tree model (Duffy (2016), Lucas (1978)). Miller and Rholes (2020) find that allowing for joint decisions, rather than individual decisions, improves choices in the optimization task (relative to the representative benchmark) by about 40%.

The model we test here relies crucially on households making optimal consumption and borrowing decisions over a life cycle. Inefficient behavior in one portion of the life cycle - under- or over-borrowing while young, for example - has meaningful implications for the aggregate economy in later parts of the life cycle. The model we test assumes agents have perfect foresight and that agents always borrow and spend along optimal paths. Thus, a potential pitfall here is the inability of real people to correctly solve the life-cycle problem. This failure would drastically decrease aggregate stability and undermine equilibration. Importantly, it could potentially mute the effects of policy interventions that operate on the economy my moderating intertemporal choice (i.e. shifting demand by manipulating rates or influencing inflation expectations).

Production economy experiments involve studying the simultaneous coordination of decisions of multiple agents in settings with input and output markets. Such experiments have been used to study patterns of international trade, exchange rates, economic growth and rationing (Lei and Noussair, 2002; Noussair et al. 1995, 1997, 2007; Fenig and Petersen, 2017). Our experiment will introduce an overlapping generations (OLG) structure to this production framework. Overlapping generations models have been taken to the lab to study inflation (Arifovic, 1995), fiscal policy (Van der Heijden et al., 1998), and coordination (Offerman et al. 2001). We make two methodological contributions to this literature. First, we introducing a novel approach to modeling OLG economies in the laboratory **(discussed in detail in the design section).** Second, we introduce a pricing algorithm based on numerical methods that facilitates laboratory experimentation with models whenever closed-form price equations are not feasible.

Expectations and Learning to Forecast In addition to the importance of optimal consumption/borrowing behavior in models of secular stagnation, there exists much experimental work to guide our implementation and understanding of expectations elicitation. Our experimental framework assume agents are able to both forecast and optimize simultaneously. In particular, subjects are able, while young, to perfectly forecast future (middle-aged) income and then borrow optimally against this forecast. This use of rational expectations to close self-referential models allows market clearing where otherwise yet unrealized information would serve as an obvious impediment. As noted by Duffy (2016), this assumption disallows testing of theoretical models using real-world empirical data since any failure of the data to match theoretical predictions could be driven either by faulty expectational assumptions or by some other aspect of the model. Much of the existing experimental literature fails to provide support for the rational expectations assumption (discussed by Camerer (1998), Ochs (1995), Duffy (2016)). Nevertheless, rational expectations dominates macroeconomic literature and so eliciting expectations is a fundamental component of testing macroeconomic theory in the laboratory.

The primary approach used within the literature is the 'learning-to-forecast' model initiated by Marimon and Sunder (1993, 1994, 1995) and developed by Heemeijer et al. (2009) and Hommes (2011), whereby subjects' elicited price forecasts simultaneously determine automated demand decisions of traders and aggregate price outcomes. This method complements the 'learning-to-optimize' approach used in the literature described earlier in this section. The combination of price expectations and decisions have been studied in Bao et al. (2012), Petersen and Winn (2014), and Petersen (2016). We incorporate both approaches into our experiment. Subjects in our experiment provide price expectations that then determine projected borrowing constraints and budgets. Subjects use these budgets to make borrowing and savings decisions.

Monetary Policy Experiments This paper also contributes to an emerging literature on liquidity traps. Arifovic and Petersen (2017) show in a learning-to-forecast experiment that expectations significantly overreact to exogenous demand shocks and that neither qualitative nor quantitative communication of a central bank's higher inflation target can effectively rescue an economy mired in a deflationary trap. However, the authors do find that quickly applied and certain fiscal stimulus can stabilize expectations and facilitate economic recovery. Hommes, Massaro, and Salle (2015) also find that fiscal stimulus is effective at mitigating deflationary spirals. Ahrens, Lustenhouwer, and Tettamanzi (2017) extend Arifovic and Petersen by allowing for human vs. robotic central bankers. They find that human central bankers can more effectively build up credibility by slowly adjusting their inflation projections upward while at the ZLB. Our paper is the first to explicitly study the effects of raising the inflation target on both expectations and real decisions.

More generally, learning-to-forecast experiments have been used to study a host of questions related to central bank communication: the construction and communication of interest rate projections and macroeconomic projections (Kryvtsov and Petersen, 2015, 2020; Mokhtarzadeh and Petersen, 2017; Ahrens et al., 2017), the communication of inflation targets (Cornand and M'Baye, 2016), and macroeconomic literacy training (Mirdamadi and Petersen, 2018). In general, these experiments highlight the importance of simple, easy-to-understand communications and information in guiding expectations to the rational expectations equilibrium.

Macroeconomists have also studied monetary policy-relevant questions in production economy settings. Bosch-Domènech and Silvestre (1998) and Lian and Plott (1997) Fenig et al. (2018) study monetary policy in a production economy where households may also trade in speculative asset markets. They show that a 'leaning-against-the wind' monetary policy that raises the nominal interest rates with asset price inflation can be effective in stabilizing asset prices with little consequence for the real side of the economy. In their environment, nominal interest rates are unbounded and can become significantly negative. This serves to fuel asset price bubbles.

Our paper builds on this research by providing the first study to compare the effects of bounded and negative interest rates on consumption-saving decisions. We find that negative interest rates propel consumers to shift consumption to the present.

2.3 Theoretical Framework and Hypotheses

2.3.1 Households

Consider an economy with young, middle-aged, and old households. Households derive utility from a single consumption good, C_t. Young households receive no income. Instead, they borrow from middle-aged households but face an exogenous borrowing constraint D_t[2], which constitutes some proportion of middle-aged income. Middle-aged households earn income from the inelastic provision of labor, $\overline{L}$, and from firm profits Z_t. These households repay debt accrued while young, and then split remaining money between consumption and savings. Old households consume with savings. A one-period, risk-free bond facilitates lending and borrowing (between middle-aged and young households) in the loanable funds market.[3] Additionally, households trade one-period debt denominated in money. The central bank controls the per-period nominal rate of return, i_t, on this asset. Thus, households maximize:

$$E_t\{ln(C_t^y) + \beta ln(C_{t+1}^m) + \beta^2(C_{t+2}^o)\} \tag{2.1}$$

subject to the following budget constraints:

[2]This is taken within the theory to be some exogenous debt constraint that the authors use to introduce debt deleveraging shocks.

[3]This implies that equilibrium in the bond market requires that the borrowing of the young match the savings of the middle-aged.

$$(1 + g_t)B_t^y = -B_t^m \tag{2.2}$$

$$C_t^y = B_t^y = \frac{D_t}{1 + r_t} \tag{2.3}$$

$$C_{t+1}^m = \frac{W_{t+1}}{P_{t+1}} L_{t+1} + \frac{Z_{t+1}}{P_{t+1}} - (1 + r_t)B_t^y + B_{t+1}^m \tag{2.4}$$

$$C_{t+2}^o = (1 + r_{t+1})B_{t+1}^m \tag{2.5}$$

$$i_t \geq 0 \forall t \tag{2.6}$$

Note that $W_t, \ P_t, \ B_t^y, \ B_t^m$ represent nominal wages, the aggregate price level, borrowing of the young, and savings for the middle-aged in some period t. Equation (2.3) implies that D_t is always binding, Equation (2.5) implies that old households consume all income, and Equation (2.6) implies a binding ZLB. This maximization problem yields the Euler equation

$$\frac{1}{C_t^m} = \beta \mathbb{E}_t \frac{1}{C_{t+1}^o} (1 + i_t) \frac{P_t}{P_{t+1}} \tag{2.7}$$

2.3.2 Firms

Firms are perfectly competitive price takers with technology $Y_t = L_t^\alpha$ that maximize profits via an optimal hiring decision:

$$\frac{W_t}{P_t} = \alpha L_t^{\alpha 1} \tag{2.8}$$

Without no source of market friction, dynamics would match those of an endowment economy and Equation (2.8) would pin down real wages. Thus, the model includes wage rigidity. Workers do not work for wages that fall below some threshold comprising a convex combination of a the flexible wage, $W^{flex} = \alpha P_t L_t^{\alpha - 1}$, and wages from the previous period, W_{t-1}. Thus, wages in this model are given by

$$W_t = max\{W_t, W_{t-1} + (1-\gamma)W^{flex}\} \tag{2.9}$$

where $\gamma \in [0,1]$ represents the degree of nominal wage rigidity in the economy. Note that $\gamma = 0, \gamma = 1$ describe fully flexible wages and complete wage rigidity, respectively. Whenever W_t equals the fully flexible wage, then an economy is experiencing inflation and labor markets clear without rationing. Otherwise, an economy experiences deflation and firms ration labor uniformly.

2.3.3 Central Bank

A mechanistic central bank sets nominal rates according to a Taylor-type monetary policy rule

$$1 + i_t = max\left(1, (1+i^*)\left(\frac{\Pi_t}{\Pi^*}\right)^{\phi_\pi}\right) \tag{2.10}$$

where i^* is the steady-state nominal interest rate, Π^* is the central bank's gross inflation target, and $\phi_\pi > 1$ is the central bank's reaction coefficient to deviations of inflation from the inflation target. Gross inflation is given by $\Pi_t = \frac{P_{t+1}}{P_t}$. The assumption of perfect foresight, combined Equation (2.7), implies a standard Fisher equation:

$$(1 + i_t) = (1 + r_t)\frac{P_t}{P_{t+1}} \tag{2.11}$$

This, coupled with the binding ZLB, places a lower bound on the inflation rate for a constant-inflation equilibrium:

$$\Pi(1 + r) = 1 + i \geq 1 \implies \overline{\Pi} \geq \frac{1}{1+r} \tag{2.12}$$

EMR note that this bound is meaningful whenever there is a permanently negative real rate because it implies that steady state inflation must remain positive.[4]

2.3.4 Equilibrium

Whenever economies face deflation, the long-term nominal rigidities prevent the emergence of market clearing wages. Thus, the real cost of wages rises causing firms to ration labor demand and an output gap to emerge. Whenever $\Pi \geq 1$, then $L_t = \overline{L}^{\alpha}, W_t = W^{flex}$, and there is no output gap so that $Y_t = \overline{L}^{\alpha} = Y^f$.

AS is then split into:

$$Y_t = \overline{L}^{\alpha}, \qquad \Pi \geq 1 \tag{2.13}$$

$$Y_t = Y_f \left(\frac{\gamma - \Pi}{\Pi(\gamma - 1)} \right)^{\frac{\alpha}{1-\alpha}}, \qquad \Pi < 1 \tag{2.14}$$

where Equation (2.13) describes the vertical portion of the AS curve and Equation (2.14) the upward sloping portion of the AS curve.[5] Note that the degree of wage rigidity γ dictates the slope of the upward sloping portion of the AS curve and can, under certain conditions, determine both the existence and uniqueness of a deflationary equilibrium.

AD is split into:

$$Y = D + \frac{(1+\beta)(1+g)D}{\beta} \frac{1}{\Pi^{\phi_\pi - 1}} \frac{(\Pi^*)^{\phi_\pi}}{(1+i*)}, \qquad i > 0 \tag{2.15}$$

$$Y = D + \frac{(1+\beta)(1+g)D}{\beta} \Pi, \qquad i = 0 \tag{2.16}$$

[4]For example, a natural rate of -4% implies a lower bound on inflation of 4%, which precludes the existence of an inflation target below 4% whenever prices are flexible. Otherwise, there would not exist an equilibrium. Hence the introduction of nominal wage rigidities as a source of market friction.

[5]The theory assumes that expectations adjust so that there is no trade off between unemployment and inflation in environments characterized by permanently high inflation. However, this same assumption does not carry over to environments characterized by low inflation or deflation. Hence, the upward sloping portion of this supply curve.

2.3.5 Monetary Policy

Using the equilibrium equations described above, we assume the following parameter values for discussion in this section: $\Pi^* = 1.1, \phi_\pi = 2, \gamma = .3, Y_f = 1, \alpha = .7, \beta = 1, g = 0, L = 1$, which are the inflation target, a measure of the central bank's responsiveness to the inflation gap, the degree of wage rigidity, the level of full-employment output, output elasticity of labor, the discount factor, population growth rate, and the inelastic labor supply, respectively. Households face two decisions during the full lifecycle: a borrowing decision while young and a consumption/savings decision while middle-aged. These two decisions drive the dynamics of this model. The concavity of the utility function is such that an optimizing middle-aged agent must consumption smooth. Savings from this cohort flows into the loanable funds market. Demand for loanable funds originates exclusively from the constrained borrowing decisions of the young.

To understand the impact of a deleveraging shock, suppose an inflationary economy faces a reduction in the amount of money that Young households may borrow for consumption. This causes a sharp decrease in the demand for loans but not the supply of loanable funds which, in turn, causes the market clearing interest rate to fall. Thus, the young who face a deleveraging shock in period t will have excess resources in period $t + 1$. This causes an increase in the supply of loanable funds in $t + 1$, further decreasing the interest rate. This translates into a decrease in AD as consumption becomes increasing desirable relative to saving. We illustrate the impact of such a shock on inflation-output dynamics in Figure 2.1.

The downward sloping portion of AD is described by Equation (2.15) and the upward sloping portion by Equation (2.16).[6] Importantly, shocks to D can eliminate a unique inflationary equilibrium and create instead a deflationary equilibrium. We set the pre-shock value of $D = 35\%$, which yields a unique inflationary equilibrium with 10% inflation. This equilibrium occurs where AD3 intersects AS in Figure 2.1. A deleveraging shock reduces the borrowing constraint to $D = 30\%$, which shifts the upward sloping demand curve in Figure 2.1 inward and yields a unique secular

[6] A shock to D impacts both segments of the curve but that this shock is offset by a simultaneous adjustment of i* in Equation (2.15) so that only the upward sloping portion of the AD curve shifts significantly upon impact of a deleveraging shock.

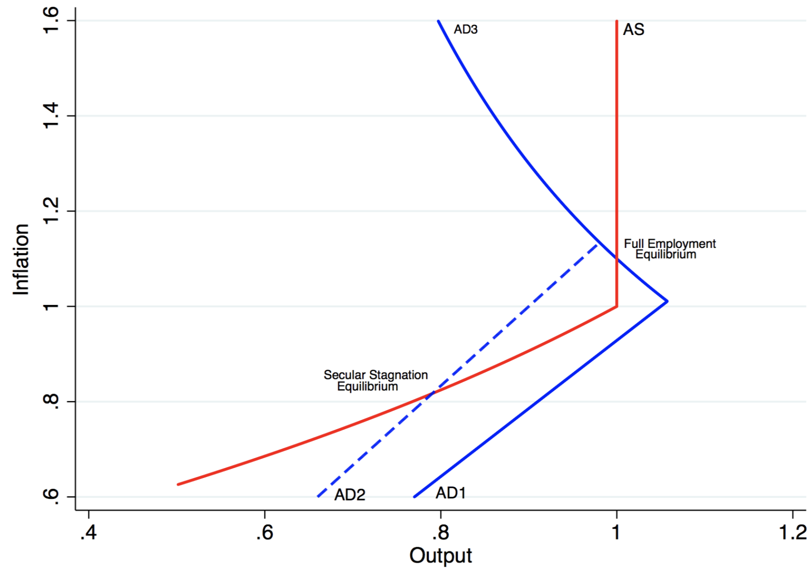

Figure 2.1: A deleveraging shock moves an economy from a unique inflationary to a unique secular stagnation equilibrium.

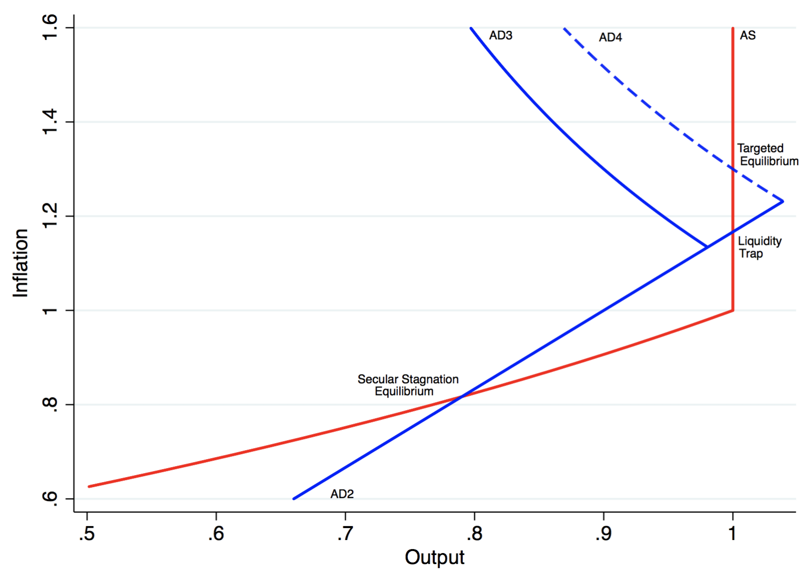

Figure 2.2: Increasing the inflation target generates two inflationary equilibria: a target, full-employment equilibrium and a liquidity trap equilibrium.

stagnation equilibrium where AD2 intersects AS. The predictions of the model yields the following two hypotheses:

H1. The economy stabilizes at the unique inflationary equilibrium in the pre-shock phase.

H2. A sufficiently large deleveraging shock will cause an economy to stabilize at the secular stagnation equilibrium.

The mechanistic central bank is bound by the ZLB and so instead addresses secular stagnation via inflation targeting. Suppose that the central bank raises its inflation target from its baseline level of to $\Pi^* = 10\%$ to $\Pi^* = 30\%$. This shifts the downward sloping sloping portion of the AD curve rightward in Figure 2.2 from AD3 to AD4, guaranteeing the existence of two inflationary equilibria: a full-employment, targeted equilibrium (where AD4 intersects the vertical portion AS) and the liquidity trap equilibrium (where AD2 intersects the vertical portion of AS). However, equilibration following this policy action hinges on the assumption that households are rational, which is a particularly vulnerable assumption.[7]

H3. Raising the inflation target to a sufficiently high level will move an economy out of secular stagnation stagnation to the targeted inflationary equilibrium.

We also consider the possibility that our mechanistic central bank can allow nominal interest rates to decrease below zero. The idea of using negative rates has gained in popularity over that last two decades as many advanced- and emerging-economy central banks find themselves constrained by the ZLB. There is some evidence that banks may be able to succesfully employ negative rates. Eggertson, Juelsrud, Summers, and Wold (2019) show that, under some conditions, negative nominal rates are expansionary. Altavilla, Burlon, Giannetti, and Holton (2019) use Euro-area data to show empirically that the transmission mechanism of monetary policy does not break down when rates become negative.

Clearly, allowing our mechanistic central bank to set negative nominal rates violates Equa-

[7] The increase in inflation target from 10 to 30% is certainly extreme, but is necessary given our parameterization that allows for more linearity in the households' utility function. Nonetheless, in thinking about future policy actions at the 2019 Jackson Hole Symposium, former Federal Reserve governor Randall Kroszner said that central bankers were searching for a "shock and awe strategy...to make sure that markets realise they're serious, and that they are going to have an impact".

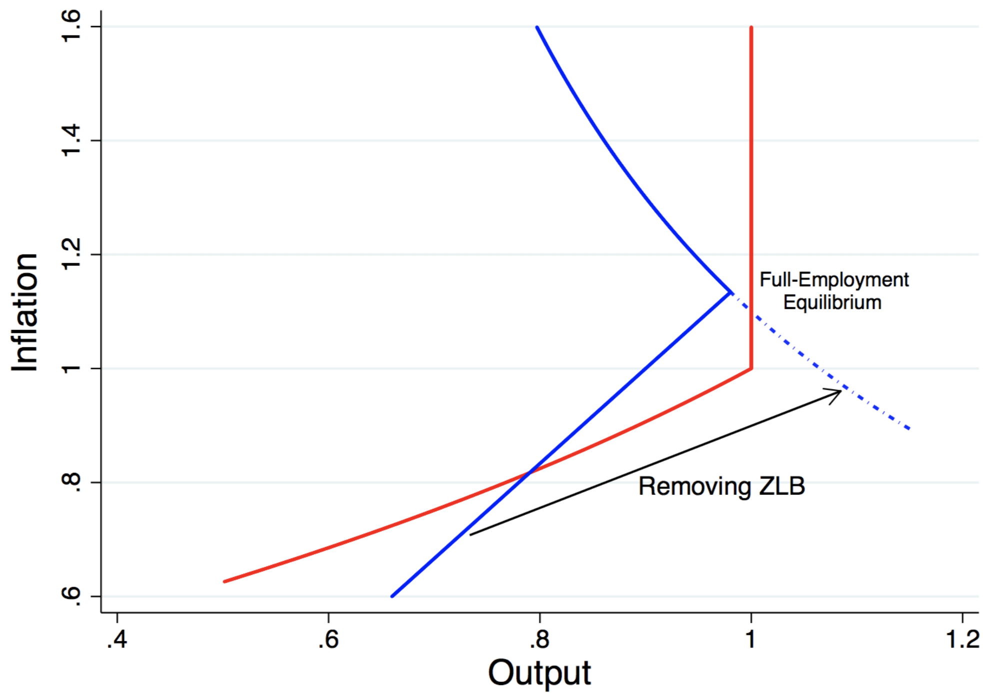

Figure 2.3: Removing the ZLB removes the kink from the AD curve, thereby eliminating the secular stagnation equilibrium and creating instead a unique inflationary equilibrium.

tion (2.6). Removing this constraint removes the kink in the AD curve, so that AD is fully describe by Equation (2.15). We show in Figure 2.3 that this change eliminates the unique secular stagnation equilibrium and recreates the full-employment equilibrium that coincides with the central bank's 10% inflation target. This yields the following hypothesis:

H4. Eliminating the ZLB moves an economy out of secular stagnation and back to the targeted inflationary equilibrium.

It is unclear ex ante whether or not this should be effective. It is quite possible that such an intervention will prove ineffective if subjects fail to fully appreciate the implications of negative interest rates. Subjects' consumption decisions may instead overreact to such an intervention, driving an economy out of a deflationary trap to some point beyond the intended inflationary equilibrium.

2.4 Experimental Implementation

We develop a new experimental environment to study expectations and economic decisions at the zero lower bound. While previous experimental work on the zero lower bound focuses solely on expectation formation, our framework will make two important contributions. First, we introduce an unexplored secular stagnation equilibrium, expanding the set of possible research questions. Second, our environment allows for the simultaneous study of expectations and financial decisions, which yields a richer and more realistic environment within which researches can explore of how policy impacts individual and aggregate outcomes. Heterogeneity in preferences toward debt and in dynamic optimization may generate important implications for the distribution of wealth that impact the efficacy of policy.

Laboratory experiments have the benefit of providing a controlled environment where researchers can clearly observe the causal effects of new and untried policies without having potentially detrimental real-world consequences. Unlike in the real world where we have only one long history, the same debt-deleveraging and policy response scenarios can be replicated with independent groups to evaluate the robustness of both behavior and policy. Compared to theory where agents' behavior is assumed, laboratory experiments allow subjects to bring their own home-grown preferences and biases into the environment.

Laboratory experiments often come with an important trade off: a sacrifice of some external validity in exchange for a gain in experimental control. As a matter of clearly discussing potential limitations of our research, we briefly discuss here what we believe are the two important potential threats to external validity: we must select a data-generating process (DGP) for our underlying economy and a subject pool to participate in our experiment. It is possible that our DGP does not map cleanly onto reality. It is also possible that our subjects do not behave in ways representative of typical economic agents.

Given our research question and design, we believe this trade off is minimal and worthwhile. Our experiments will involve financially-incentivized undergraduate students in the role of households making financial decisions for multiple generations. Undergraduate students in laboratory experiment have been shown to behave similarly with real world agents in markets, financial decision making, and in expectation formation (Kessler and Vesterlund, 2015, Cornand and Hubert, 2018). The data-generating process will be based on EMR's model, which has been a major advancement in capturing realistic inflation dynamics. That is, this design represents, arguably, a more realistic economic environment than, for example, environments where interest rates cannot remain permanently negative due to a representative discount factor. Our economic environment subsumes the inflationary and liquidity trap equilibria present in such models, while also allowing for a secular stagnation equilibrium to emerge.

2.4.1 Experiment

There are several goals for this experiment. First, we seek to establish whether secular stagnation equilibria emerge in an experimental economy driven by the combination of exogenous demand shocks and the real decisions of economic agents. Second, we test whether unconventional monetary policies such as raising the inflation target or eliminating the ZLB can effectively stimulate inflationary expectations and aggregate demand at the ZLB. Third, we test whether these unconventional policy can rescue economies from secular stagnation.

Each laboratory session - which is an independent economy - consists of 21 household agents that form price forecasts and budgetary decisions in a three-period overlapping generations framework. For simplicity, our design involves the automation of young and old households. The young decisions are completely automated to ensure that the young are maximally leveraged.[8] Each economy has seven undergraduate participants that make repeated budget decisions that influence the consumption outcomes of the 14 remaining middle-aged and old households.

The experiment consists of 50 periods (30 periods in our baseline treatment where an intervention is not introduced). Each period t consists of three stages.

Stage 1: All subjects simultaneously submit forecasts about the current price, $E_{i,t}P_t$ and sub-

[8]Small deviations in borrowing can produce drastically different inflation and output dynamics. In pilot experiments that did not involve automated young households, we found that subjects consistently underborrow during the early periods of their lifecycle (consistent with debt aversion observed by Meissner, 2016) which yields unstable, deflationary economies. We therefore automate this decision so that we create an experimental environment that allows us to test policy prescriptions aimed at alleviating secular stagnation etc.

sequent period price, $E_{i,t}P_{t+1}$. Subjects also submit a qualitative forecast about the change in the nominal interest rate relative to the previous period (increase, stay the same, decrease). Subjects earn 2 points for correct qualitative interest rate forecasts and zero points otherwise. We incentivize price forecasts using the following payoff function:

$$ForecastPoints_t = 2^{-|E_t P_t - P_t|} + 2^{-|E_{t-1} P_t - P_t|} \tag{2.17}$$

Note that subjects can earn a maximum of 4 points per period for perfect price forecasts. Forecasting points for either forecast drop by one half for each lab dollar that a subject under or over forecasts. This is depicted graphically in Figure 2.4

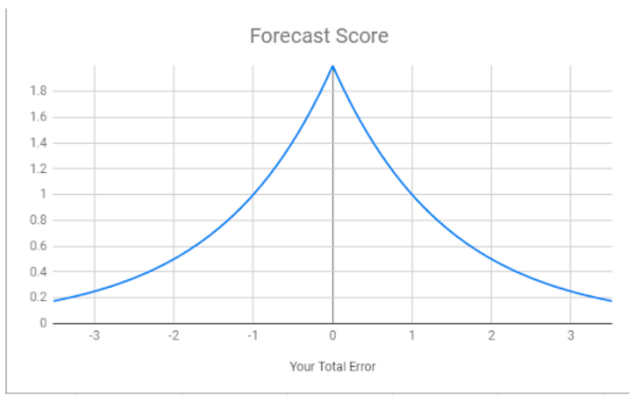

Figure 2.4: Payoff function for an individual price forecast

We elicit price and interest rate forecasts for various reasons. The efficacy of our policy interventions hinges critically on whether participants form model-consistent expectations. Price forecasts provide valuable insight into the relative importance of expectations in informing spending decisions. The median concurrent price forecast is also used to compute the expected income and interest rates participants are likely to face in the current period. We present this information to participants in the subsequent phase to facilitate their spending decisions.[9] While the qualitative interest rate forecast has no impact on how the economy evolves, it allow us to assess participants' comprehension of the central bank's policy rule.

Stage 2: Participants play the role of a middle-aged household who makes a consumption-saving decision after receiving predictions about their concurrent income, nominal interest rate, and concurrent and subsequent prices. Specifically, each participant makes a decision about how much of their current nominal income, $E_t Y_{i,t}$ to spend in period t. Any unspent income is automatically saved and consumed in the subsequent period when the agent becomes old.

Participants earn points based on their consumption decisions while middle-aged and old. We induce participants to behave as if they had a per-period utility function given by $U_t = 5 + ln(\alpha + C_t)$ where $\alpha = 0.00673$.[10] Since a participant is essentially two households at any point in time (the middle-aged agent they are actively making a decision for in period t and the old-agent that they previously made a decision for in period $t-1$), they possess two separate utility functions.

Importantly, any money saved by a middle-aged household in period t will accrue interest at the prevailing nominal interest rate and be available to the same household in period $t+1$. This now-old household will spend all remaining wealth on output and the subject receives all utility from this consumption of output. The timing structure is depicted graphically in Figure 3.2.

After all participants have submitted their spending decision, we use all automated young agents and middle-aged participants' spending decisions in period t, as well as the remaining spending dollars of the period t old agents determined in period $t-1$, to compute total period t dollars for consumption spending. This information is used to clear markets, allocate output and assign utility. Subjects earn points based on how much of the consumption good (output) they purchase, on the accuracy of their price forecasts, and on the accuracy of their qualitative interest rate forecast.

[9]We use a median price prediction rather than a weighted average price. This prevents subjects from coordinating expectations to manipulate prices in any favorable way. This also mitigates the ability of any one subject to significantly influence the economy via extreme expectations.

[10]This is a slight modification of the utility function to allow for the possibility that subjects consume zero output without facing exponentially negative payoffs. We select α such that $U(C_t) = 0$ when $C_t = 0$.

Period	Middle-Aged Household	Old Household (Automated)
1	Decision 1 (e.g. spend 40% / saving 60%)	
2	Decision 2	Decision 1 (e.g. spend remaining 60% + interest earned)
3	Decision 3	Decision 2
4	Decision 4	Decision 3
5	...	Decision 4

Figure 2.5: Co-determination of middle-aged and old spending

Stage 3: The third phase is for participants to review the outcomes of the current period. All participants observe the total amount of output produced, price of output, nominal interest rate, as well as their own current spending decision and the amount of points earned from consumption.

We provide a history of all aggregate-level variables to all subjects in all periods (following the first period) during both stages of each period. Additionally, the central bank informs all subjects of its policy rule and its current inflation target during both stages of each period.

We convert experimental points into real dollars at a rate of 20-to-1. The use of monetary incentives in the experiment reflects the position of the economics profession that experimental studies generalize to situations outside of the laboratory only if subjects' decisions have a direct and significant influence on their compensation.

We provide subjects with two tools to facilitate play. The first tool, available in Stage 1 of each period, allows subjects to convert inflation expectations into price expectations or to convert price expectations into inflation expectations. We do this so that subjects can easily incorporate both inflation and price information when forming price forecasts. The second tool, available in Stage 2, takes as inputs a subjects' price expectations and returns to them a suggested level of spending conditional on their individual price expectations. We note to subjects in our instructions that this suggested level of spending is conditional on their expectations and also inform them that they may enter any strictly positive number for their expectations. They are beholden to neither the price prediction provided in Stage 1 nor the median price predictions displayed to them on the

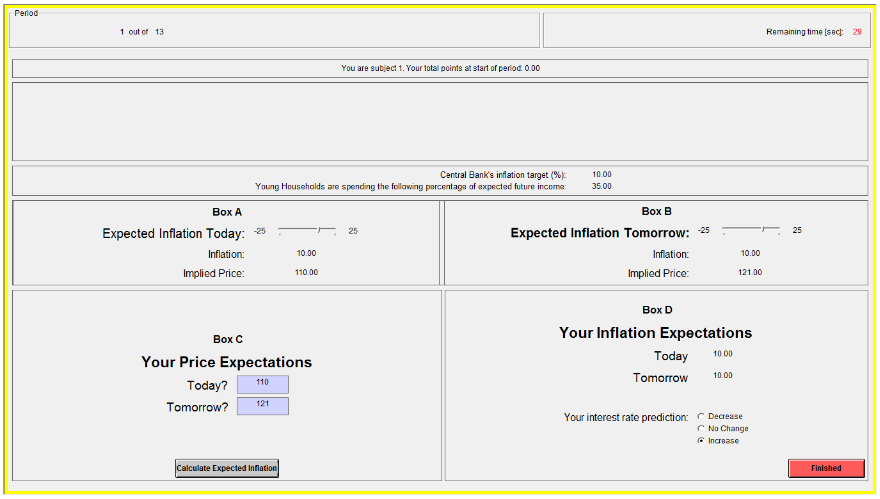

Figure 2.6: Stage 1 Screen

Stage 2 screen. Finally, we also provide subjects with a full history of aggregate outcomes and individual decisions on all screens of the game. For examples, refer to the screen shots depicted in Figure 2.6 and Figure 2.7.

We face the non-trivial challenge of simultaneously clearing markets and allowing young agents to borrow from future uncertain income. EMR assume rational expectations to assuage these thorny issues. We follow their approach by relying on subject-provided expectations to provide income, price, and interest rate signals which can inform participants' decisions before markets are cleared. We couple subjects' expectations with a novel pricing algorithm to determine aggregate spending, price, wage, output, labor demand, and interest rate.

We develop and implement an algorithm built on numerical methods for solving the pricing equation. This is necessary because there is no analytical solution for prices, given the structure of this self-referential model. We solve for prices within our framework as follows:

For the program, we define the per-period, market-clearing price as:

$$P_t = \frac{C_Y + C_M + C_O}{Y_t}. \tag{2.18}$$

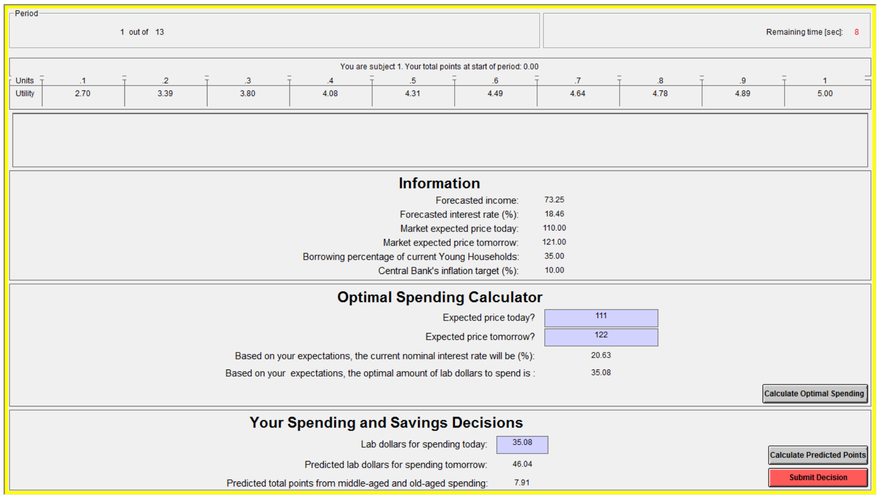

Figure 2.7: Stage 2 Screen

This gives us the following piece-wise, per-period price function:

$$P_t = \frac{C_Y + C_M + C_O}{Y^f}, \qquad \Pi \geq 1 \tag{2.19}$$

$$P_t = \frac{C_Y + C_M + C_O}{Y_t}, \qquad \Pi < 1. \tag{2.20}$$

We proceed by first supposing that prices are determined by Equation (2.20), which can be rewritten as

$$P_t = \frac{C_Y + C_M + C_O}{\frac{w_t}{P_t \alpha}^{\frac{\alpha}{\alpha - 1}}}. \tag{2.21}$$

Isolating P_t yields

$$P_t^{\frac{-1}{\alpha - 1}} = \frac{C_Y + C_M + C_O}{\frac{w_t}{\alpha}^{\frac{\alpha}{\alpha - 1}}}. \tag{2.22}$$

However, we know that per-period wages w_t are also a function of P_t since we have that

$$w_t = max(P_t\alpha\overline{L}^{\alpha-1},\ \gamma w_{t-1} + (1-\gamma)P_t\alpha\overline{L}^{\alpha-1}). \tag{2.23}$$

Substituting yields

$$P_t^{\frac{1}{\alpha}} = (C_Y + C_M + C_O)^{\frac{1-\alpha}{\alpha}} w_t\alpha^{-1} = C(\gamma(w_{t-1}) + ([1-\gamma]P_t\alpha\overline{L}^{\alpha-1}) \tag{2.24}$$

where

$$C = (C_Y + C_M + C_O)^{\frac{1-\alpha}{\alpha}}\alpha^{-1}.$$

Collecting prices, factoring, and making the following variable substitutions,

$$i.\quad b = \frac{1-\alpha}{\alpha}$$

$$ii.\quad A = C(1-\gamma)\alpha\overline{L}^{\alpha-1}$$

$$iii.\quad B = C\gamma w_{t-1}$$

,

yields

$$P_t[P_t^b - A] = B. \tag{2.25}$$

We solve this via the Newton-Raphson method of numerical approximation.

For example, suppose $f(P_t) = P_t^{b+1} - AP_t - B = 0$ and define an initial guess for our price as $X_0 = P_{t-1}$ and some stopping rule predicated upon meeting some minimum error rate ϕ. Then, if $f(X_0) \leq \phi$ the algorithm stops and $P_t \approx P_{t-1}$. Otherwise, if $f(X_0) > \phi$ the algorithm proceeds as follows:

$$X_1 = X_0 - \frac{f(X_0)}{f'(X_0)} = P_{t1}\frac{P_{t-1}^{b+1}AP_{t-1} - B}{(b+1)P_{t-1}^{b} - A}$$

Once the algorithm arrives at some X_i such that $f(X_i) \leq \phi$, define a temporary price as $P_t \approx X_i$.

Finally, we calculate output given this price and aggregate spending. If output exceeds potential then we know that our assumption that Equation (2.20) determines prices in a given period is incorrect, and the algorithm instead sets prices according to Equation (2.19).

Below we summarize how the Young, Middle-Aged, and Old households' spending decisions are computed.

The Young: We automate young households such that young household i in period t will automatically borrow a proportion $D_t^i \in [0, 1]$ of its middle-aged income. This is clearly problematic since the middle-aged income of these agents is not actually determined until markets clear in the following period. However, EMR's assumption of rational expectations circumvents this issue in theory. If subjects' expectations are rational, then there should be no difference in expected and realized prices. Thus, we compute the consumption expenditure of a Young household i in period t as $c_t^{i,y} = D_t^i E_t\{P_{t+1}\} = D_t^i E_t[Y_{t+1}^{i,m}]$.

The Middle-Aged: Subjects make budgeting decisions as Middle-aged households in stage 2 of each period. Subjects here have two considerations: a debt-repayment obligation incurred automatically by the household when young and a consumption/savings decision. However, our subjects face these considerations before income is actually determined. To deal with this, we suppose that the income of Middle-aged households is equivalent to the market expectation for current period prices. If we again suppose that we are in period $t = 0$ then middle-aged income is given as $E_0\{P_0\}$. Thus, a Middle-aged subject has net income $ni = E_0\{P_0\}c_1^j(1 + i_1)$. Suppose α_j is the proportion of net income allocated to savings so that $c_0^{j,m} = \alpha_j(E_0\{P_0\}c_1^j(1 + i_1))$. We can then use $c_0^{j,m}$ for market clearing, which finally informs us how much money Middle-aged agents actually earn. One issue here is that if subjects systematically deviate from RE then we can have that $E_0\{P_0\} > P_0$ or $E_0\{P_0\} < P_0$. Because we have already cleared markets using a consumption level based on the market-expected price, we occasionally must adjust for deviations from RE by changing the amount of money that these middle-aged agents hold in savings for consumption while old. This is because deviations from RE (and from the correspond consumption/savings decision) can drive a wedge between expected and realized savings for middle-aged agents. If we call t expected savings of middle-aged agent j as $E_0\{s_0^{m,j}\}$ then $E_0\{P_0\} > P_0$ would cause $E_{j,0}\{s_0^{j,m}\} > s_0^{m,j}$ and opposite for the opposite case.

The Old: The decision for a given old household in period t is determined by the budgeting decision of a subject acting as a middle-aged household in period $t-1$. Old households automatically spend all remaining wealth on output. For example, if a subject i assigned a Middle-aged household in period $t-1$ instructs its household to save $s_{t-1}^{m,i}$ dollars then in period t that household will allocate $p_t c_t^{t,i} = s_{t-1}^{m,i}(1 + i_{t-1})$ to consumption dollars.[11] Note then that subjects, anytime following period 1, earn consumption points from a currently-assigned Middle-aged household and an Old household. This Old household is the Middle-aged household assigned to that subject in the previous period.[12]

2.4.2 Treatments

We conduct a series of treatments to explore the learnability and stability of different equilibria with and without policy action. We initialize all sessions at the unique full-employment equilibrium where we assume, for steady state values, that the economy is operating along the steady state inflation path. A surprise exogenous deleveraging shock moves the borrowing constraint from $D = .35$ to $D^{shock} = .3$ [13] to create a unique secular stagnation equilibrium. Our interest is in the ability of different unconventional monetary policy actions to move the experimental economies out of secular stagnation and back to a full-employment equilibrium. Monetary

[11] Here, c_t represents units of output and $s_{t-1}^{m,i}$ is the dollar amount saved by agent i in period $t-1$ while middle-aged.

[12] We automate Old households in period 1 based on the assumption that the economy moved along the steady-state inflation path in all periods before the start of our experiment.

[13] In our baseline treatment (*Baseline*), which does not feature a policy intervention, we set $D^{shock} = .28$.

policy interventions include raising the inflation target and allowing a negative nominal interest rate. Raising the inflation target generates multiple equilibria: a secular stagnation equilibrium, a liquidity trap equilibrium, and the targeted, high-inflation equilibrium. Removing the ZLB eliminates the upward sloping portion of the aggregate demand curve so that AD is fully describe by Equation (2.15). In theory, allowing for negative nominal rates can fully offset the deleveraging shock and stimulate aggregate demand sufficiently to recreate a unique inflationary equilibrium that coincides with the central bank's inflation target.

We use π^{tgt}, π^{ss}, π^{lt} to denote an inflationary steady state equilibrium, a secular stagnation steady state equilibrium, and a liquidity trap equilibrium, respectively.

2.4.2.1 Baseline

Baseline explores the ability of subjects to coordinate on the unique equilibria when they exist. This treatment features 30 periods of play divided into a 15-period pre-shock phase and a 15-period post-shock phase. The pre-shock phase features a unique equilibrium of π=10% with full-employment and output. This is followed by a deleveraging shock in period 16 that moves $D_t =$ $.35$ to $D_t^{shock} = .28$, which creates a unique secular stagnation equilibrium of $\pi^{ss} = -24.4\%$ with labor rationing and output well below potential. We announce the deleveraging shock to subjects at the beginning of period 16 before the begin stage 1. This announcement informs subjects about the deleveraging shock, the magnitude of this shock, a brief explanation of how it impacts the economy, and informs subjects that the shock is permanent. Further, all subjects know that all other subjects know about the shock.

2.4.2.2 Policy treatments

Each subsequent treatment embeds *Baseline*, with the caveat that shocks in *HigherTarget* and *NegativeIR* are such that $D_t = .35$ to $D_t^{shock} = .3$.

HigherTarget This treatment features 50 periods of play divided into three phases. The first two phases are fully described by *Baseline*. The third phase begins when the central bank announces the change to its inflation target in the beginning of period 30. The higher inflation target generates an environment with three equilibria: a secular stagnation, a liquidity trap, and a new targeted inflationary equilibrium. The policy intervention is announced on screen and aloud so that all subjects know that all subjects know about the policy change.

NegativeIR This treatment again involves three phases where the first two phases are fully described by *Baseline*. Phases 3 begins following removal of the ZLB in period 30. Doing this removes the kink point in the AS curve, creating a unique, inflationary equilibrium that coincides with the central bank's inflation target. The policy intervention is announced aloud and accompanied by a single-page document that provides an example to help subjects understand the implications of negative interest rates.

We summarize the treatment parameterization in Table 2.1 below.

	Phase 1	**Phase 2**	**Phase 3**
Baseline	$\pi^{tgt} = 10\%$	$\pi^{ss} = -24.7\%$	N/A
HigherTarget	$\pi^{tgt} = 10\%$	$\pi^{ss} = -18.3\%$	$\pi^{tgt} = 30\%, \pi^{ss} = -18.3\%, \pi^{lt} = 16.7\%$
NegativeIR	$\pi^{tgt} = 10\%$	$\pi^{ss} = -18.3\%$	$\pi^{tgt} = 10\%$

Table 2.1: Parameterized equilibria across treatments and phases

2.5 Results

We begin with results from *Baseline*, where we initialize our experimental economies with a unique inflationary equilibrium and then replace this with unique secular stagnation equilibrium via a sharp deleveraging shock. Results for this treatment are shown in Figure 2.8. This figure plots individual inflation expectations, the median inflation expectation, aggregate inflation (top panel), and individual consumption demand (bottom panel) from each *Baseline* session. Note that each session-level graph is split into two phases: a pre-shock phase and a post-shock phase. The vertical, dashed line denotes the period of the deleveraging shock and the two horizontal lines denote the pre-shock, full-employment and the post-shock, secular stagnation equilibrium levels of inflation (top panel) and consumption (bottom panel).

In four of the eight sessions, the economies appear to be converging to the steady state equilibrium in the pre-shock phase of each session. This is particularly impressive given that, as we discuss later in the results section, the overwhelming majority of subjects abstain from simply forecasting the central bank's inflation target in each period and instead adapt a forecasting heuristic that involves updating as a function of recent economic outcomes. The deleveraging shock, which occurs in period 15, consistently generates deflation. However, not all economies completely converge to the secular stagnation equilibrium following this shock. Instead, we see varying degrees of deflation emerge ranging from mild deflation to severe deflation at the secular stagnation equilibrium. The inability of economies to completely converge to the secular stagnation is most often attributable to over-consumption that emerges once subjects begin to experience deflation (see Figure 2.10). That is, we find evidence of a Pigou effect (Patinkin, 1948), whereby deflationary episodes are quelled by a wealth effect generated by increasing real wealth balances.

The economies in sessions 3 and 6, however, experience rampant inflation. This is driven by a confluence of highly optimistic expectations and very large spending shocks. This is striking given that the central bank pursues an aggressive policy response to inflation and is unrestrained in adjusting its policy rate upwards. In fact, as expectations remain anchored despite increasing rates, we eventually see that pursuing a Taylor-type rule reinforces rather than dampens this inflationary pressure. The lack of responsiveness of the economy to the high interest rate suggests that the wealth effect strongly dominates the substitution effect for participants.

Worth noting in Figure 2.8 is how consumption changes in response to the deleveraging shock. Session 8 gives a nice example. Though consumption is relatively well anchored in phase 1 of each session, we see that the deleveraging shock and subsequent deflationary pressure drastically increases consumption heterogeneity, which suggests that such shocks reduce aggregate welfare but the decrease in welfare is larger for some than others. This pattern emerges in most of our experimental economies.

Figure 2.9 follows the same conventions as Figure 2.8 but plots the nominal interest rate rather than nominal inflation. We see here that the central bank is constrained by the zero lower bound

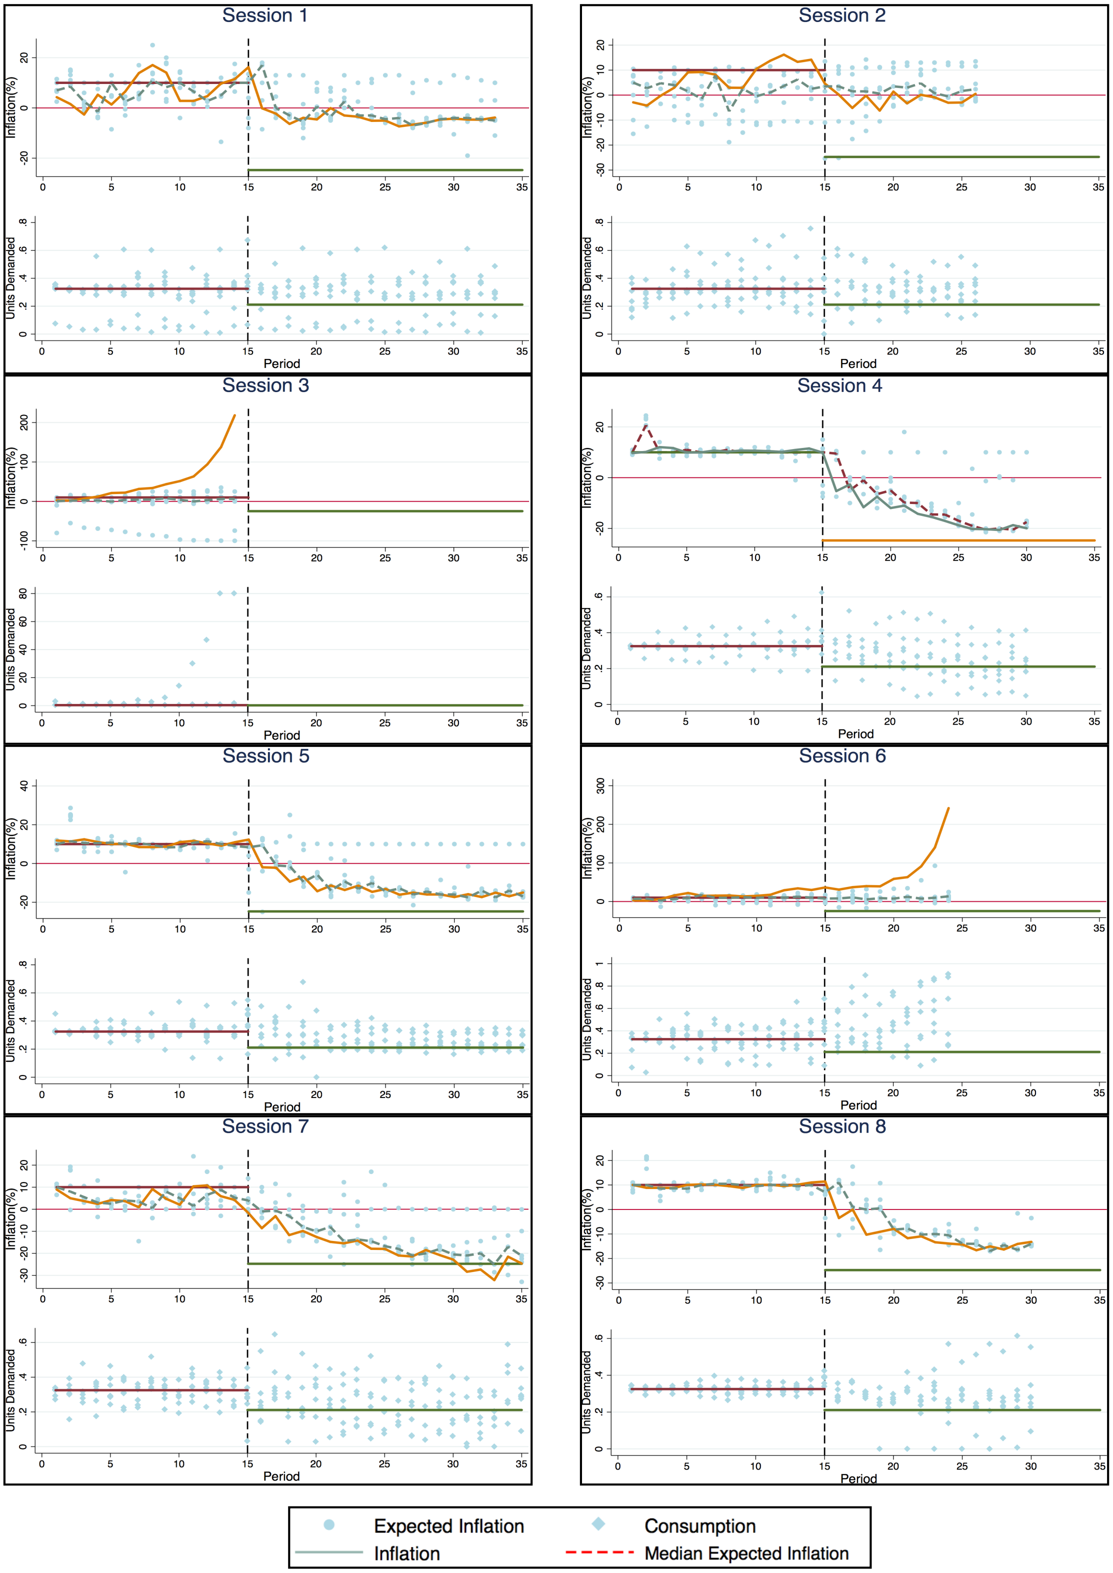

Figure 2.8: Aggregate inflation by session for *Baseline*

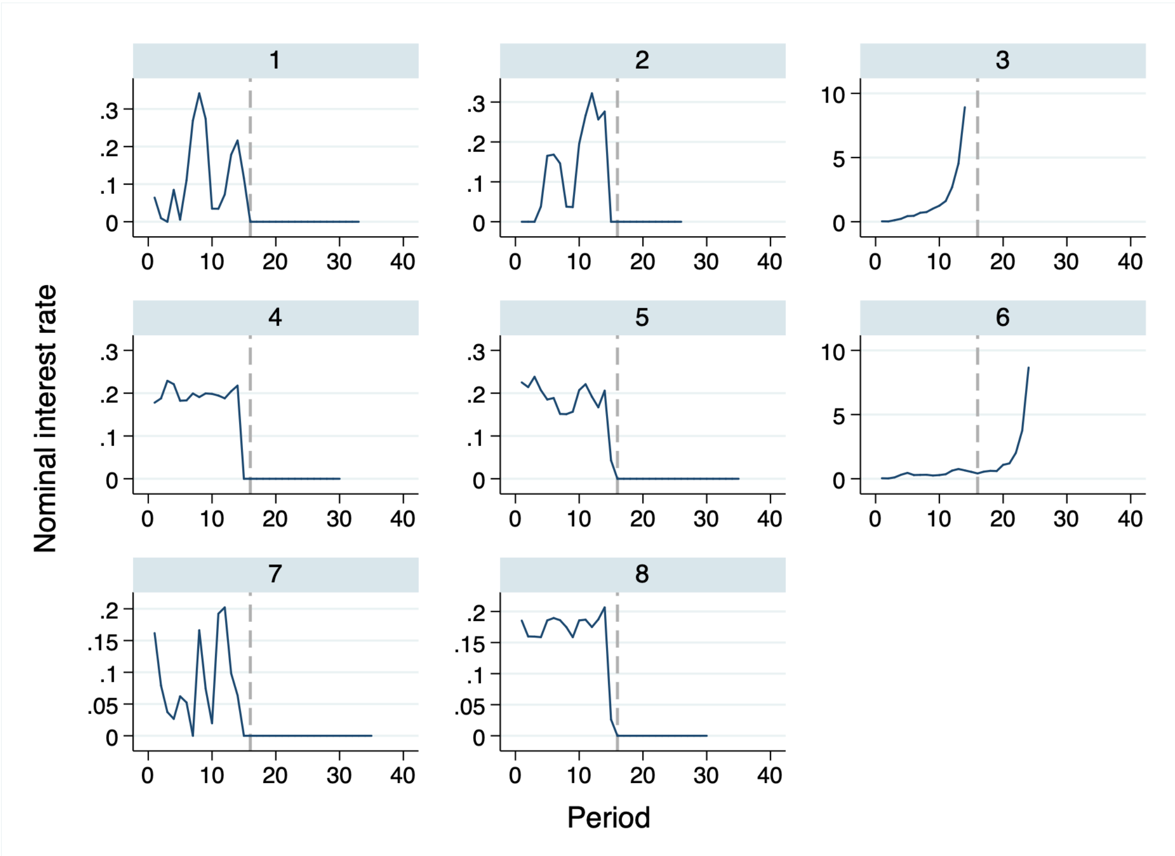

Figure 2.9: Nominal interest rates by session for *Baseline*

almost immediately following a deleveraging shock. Nevertheless, this drastic reduction in interest rates does little to stimulate aggregate demand.

Additionally, we examine consumption demand aggregated across phases one and two of *Baseline*. We show this in Figure 2.10, which plots the cumulative density function (CDF) for aggregate demand of all participants across all periods for phases one and two of *Baseline*. The two vertical lines denote optimal consumption demand for an individual under rational expectations. The leftmost vertical line denotes optimal post-shock consumption while the rightmost line denotes optimal pre-shock consumption. The dashed density line is the post-shock CDF.

There are several things worth noting in Figure 2.10. First, we see that in Phase 1 of *Baseline*, our subjects, on average, consume approximately optimally relative to a fully rational economic

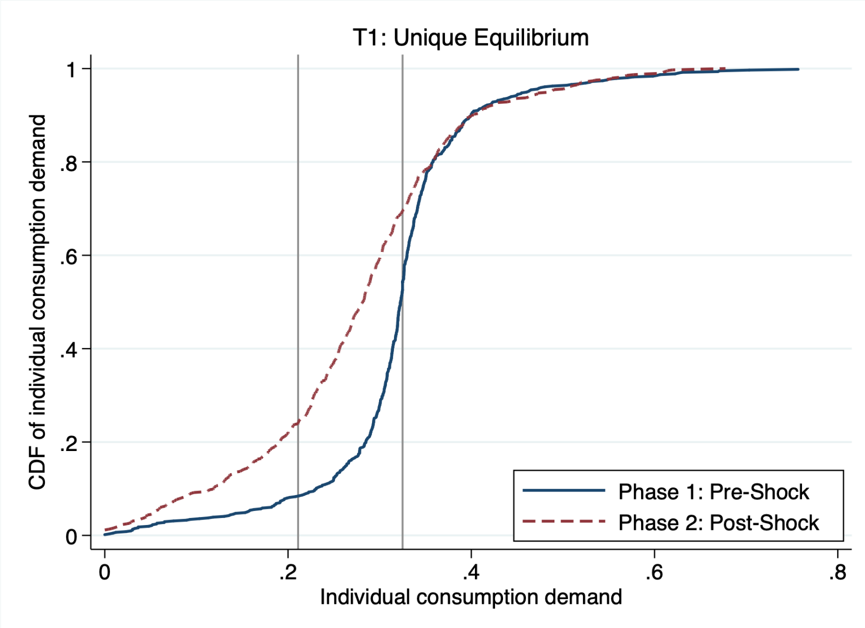

Figure 2.10: Nominal interest rates by session for *Baseline*

agent. Further, we see that post-shock consumption is much higher than that predicted, which is the predominant reason that some economies do not converge fully to the secular stagnation equilibrium. This pattern of over consumption following a deleveraging shock emerges consistently across treatments.

HigherTarget embeds *Baseline* and explores aggressive inflation targeting as a method of pulling an economy from secular stagnation. We show results from this treatment in Figure 2.11. We split *HigherTarget* session into three phases (from left to right): a pre-shock phase, a post-shock phase, and a post-intervention phase. Vertical dashed lines denote the timing of the deleveraging shock and the central bank's policy intervention, respectively. Horizontal lines denote equilibrium inflation and consumption levels. Note that phase 3 of *HigherTarget* features three horizontal lines. These correspond to the three equilibria depicted above in Figure 3. The top orange line is the full-employment equilibrium that coincides with the central bank's new inflation target, the grey dashed line (middle line in phase 3) denotes the liquidity trap equilibrium, and the solid bottom green line denotes the secular stagnation equilibrium.

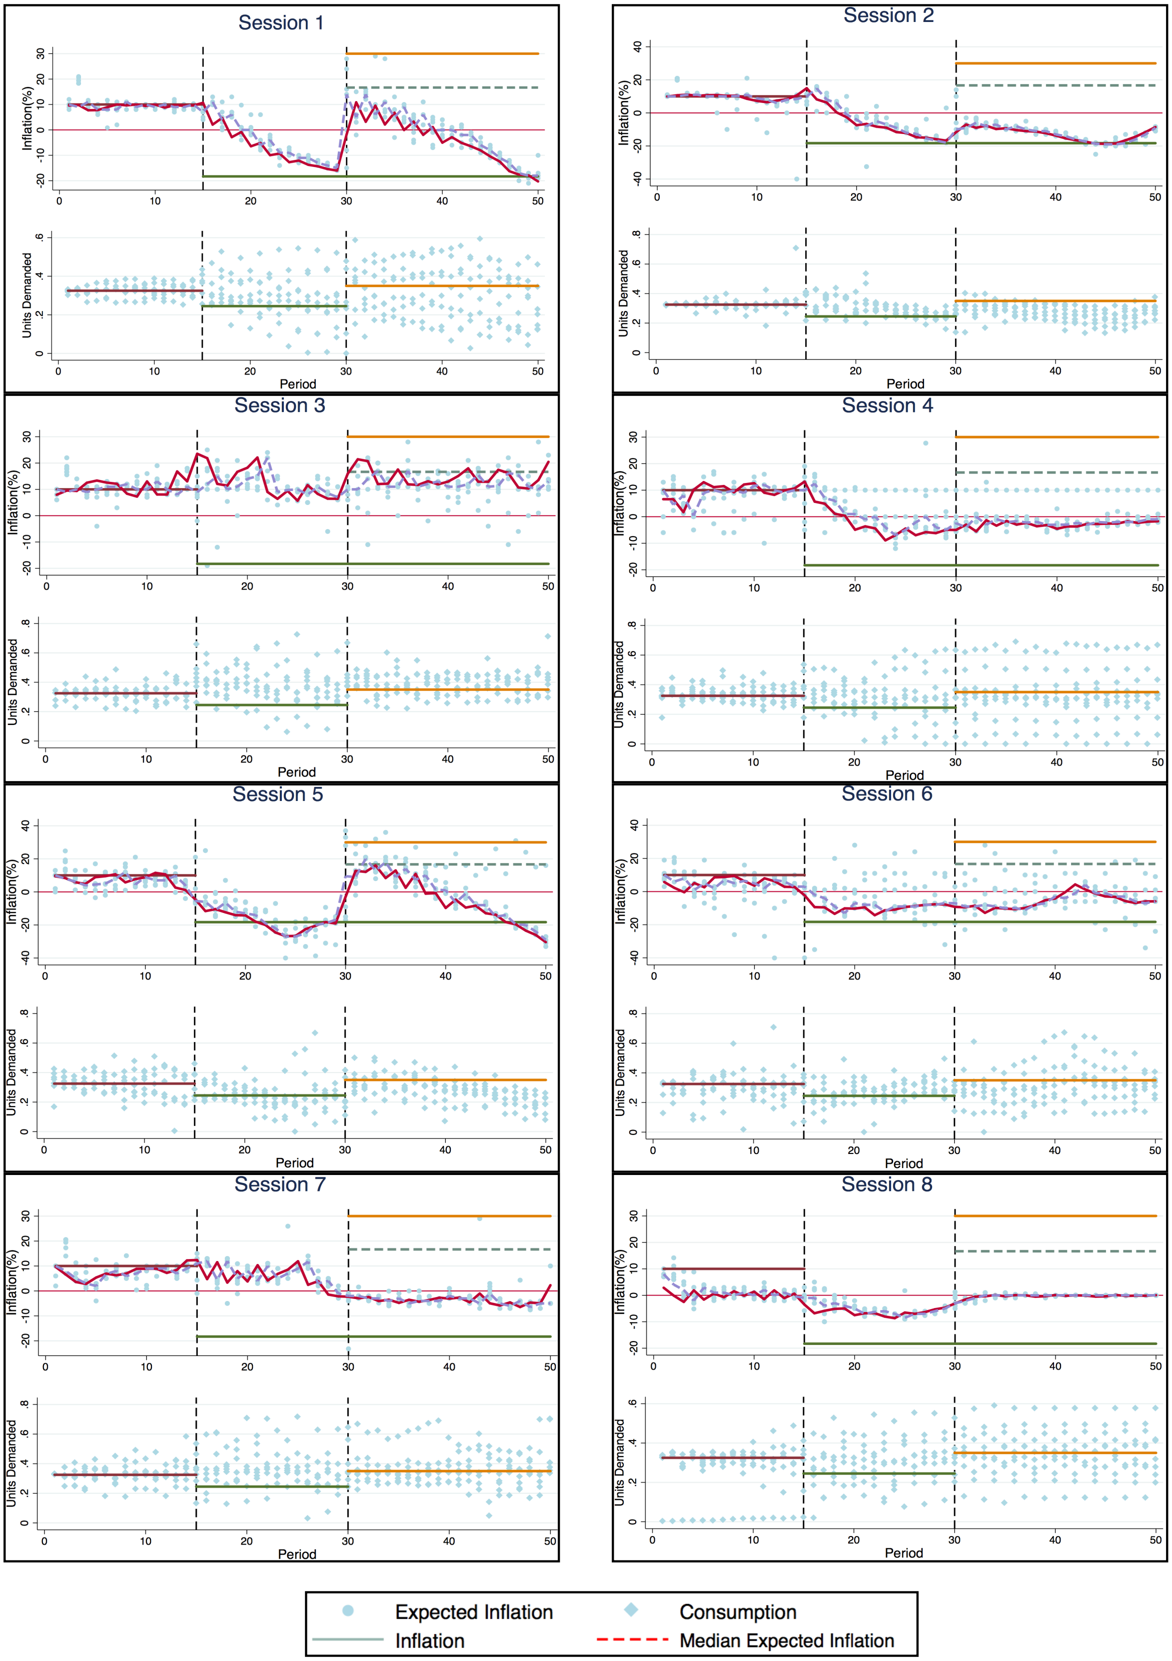

Figure 2.11: Aggregate inflation dynamics by session for *HigherTarget*

Again, subjects in this setting do a good job of playing the inflationary equilibrium and we see that the deleveraging shock generates pessimistic expectations and deflationary pressure on the individual economies. Following deflationary episodes generated by the deleveraging shock, our mechanistic central bank attempts to address secular stagnation by increasing its inflation target. The idea here is that this new inflation target will generate optimistic, inflationary expectations that will fuel spending. This expectations-driven increase in demand in turn increases the borrowing and spending capacity of the Young and Middle-aged households through increased expected nominal income.

The increase in the inflation target ubiquitously fails to return the economies to the desired, full-employment equilibrium. Instead, we observe in most economies that inflation reacts very little to the increase in the inflation target. In economies where there is some initial reaction to the announcement/target increase (sessions 1, 2, and 5), there is a slow but steady collapse in inflation expectations and inflation, and a subsequent return to the secular stagnation equilibrium.

Also noteworthy is that some economies in the post-intervention environment do manage to mitigate deflation by coordinating on zero inflation. Zero inflation often becomes a focal point for our economies, despite not being part of the set of predicted rational expectations equilibria. However, we find that the result of over-consumption that begins after subjects have experienced prolonged deflation. Coordination on zero-percent inflation has the effect of trivializing price forecasts for subjects and greatly reducing the complexity of the two-period optimization problem subjects face in stage 2 of each period.

As was true in Figure 2.8, we see a considerable increase in consumption heterogeneity following the deleveraging shock. Worth noting here is that increasing the inflation target does little to coordinate inflation expectations on the central banks target and also has little coordinating effect on consumption.

We see in Figure 2.12 that nominal rates in this economy often exhibit quite a bit of volatility pre-shock, which yields relatively stable pre-shock inflation dynamics. And as was true in *Baseline*, we see that almost all of our economies converge almost immediately to the zero lower

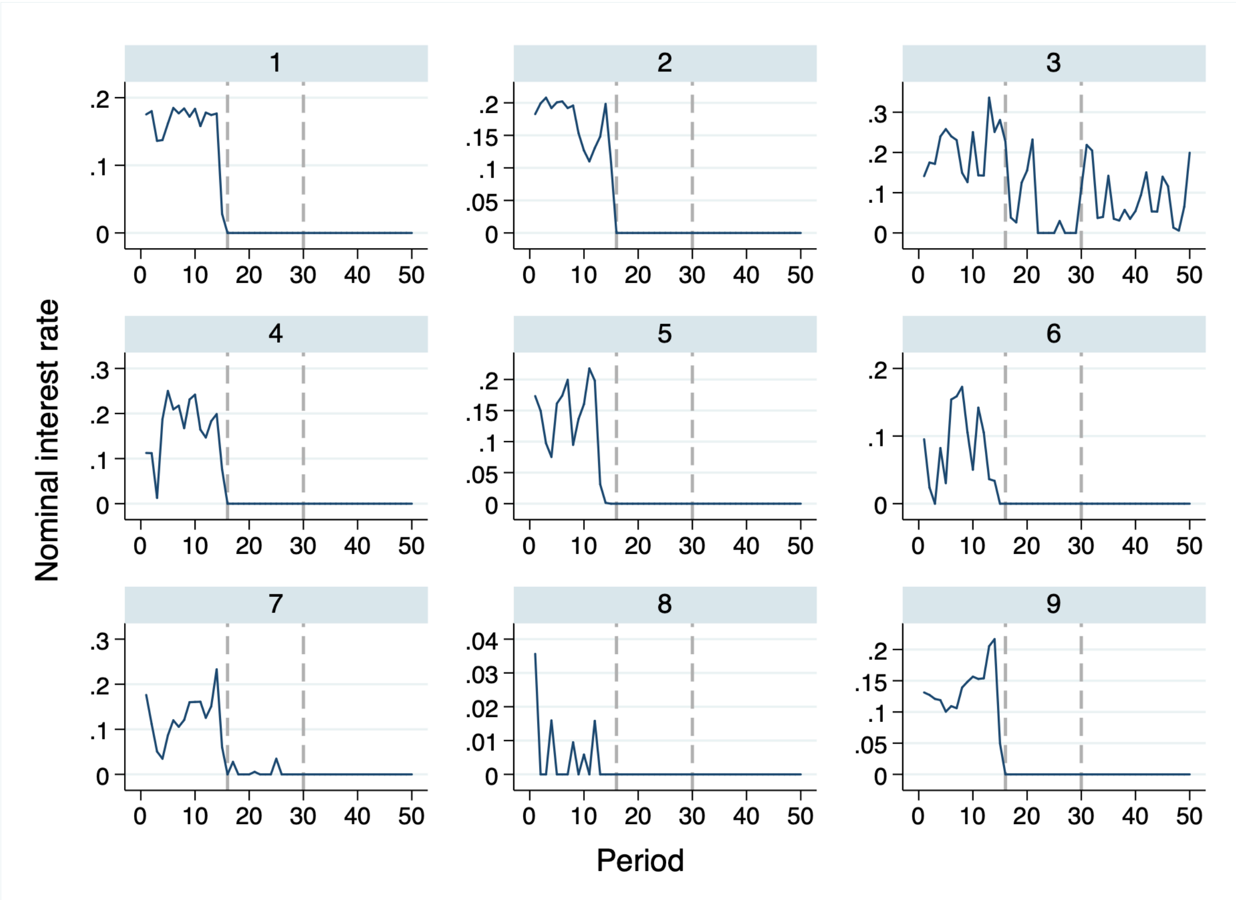

Figure 2.12: Nominal interest rates by session for *HigherTarget*

bound following our deleveraging shock. We see that from each of our economies[14], regardless of post-intervention dynamics, remains constrained at the zero lower bound post intervention.

We also consider consumption dynamics for *HigherTarget*, which we show in Figure 2.13. This graph is similar to the consumption graph from *Baseline*. However, we now display three density lines and three vertical lines depicting optimal consumption levels under rational expectations for each of the three phases of *HigherTarget*. The three vertical lines denote optimal levels (from left to right) for post-shock consumption, pre-shock consumption, and post-intervention consumption. The solid and dashed density lines depict the CDFs of phase 1 and phase 2 consumption, respectively. The dashed-and-dotted density curve depicts the CDF of post-intervention consumption.

[14] Aside from session 3.

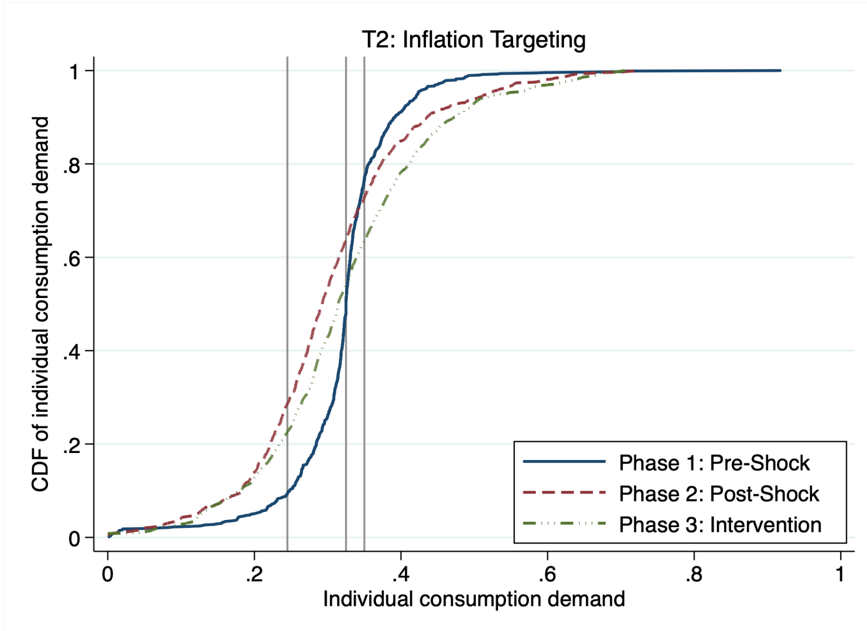

Figure 2.13: CDFs of consumption for all subjects in all periods of *HigherTarget*

As was true in *Baseline*, we see in Figure 2.13 that, on average, subjects consume optimally in Phase 1 but over consume in Phase 2. We also see in Figure 2.13 that subjects under consume in Phase 3, which is driven at least in part by the inability of the central bank to re-anchor expectations on its higher inflation target. In additional to low demand of Middle-aged driven by pessimistic expectations, this sub-optimal demand is also driven by Young households borrowing against future income that is much lower than is theoretically predicted under rational expectations.

NegativeIR also embeds *Baseline*. However, the central bank in these experimental economies now intervenes by permanently removing the ZLB in an effort to stimulate inflation expectations and increase aggregate demand in the hopes of returning the stagnating economies to the full-employment equilibrium. Results from this treatment are shown in Figure 2.14.

As in *HigherTarget*, the central bank's policy intervention occurs at the beginning of period 30. There are two key differences between *NegativeIR* and *HigherTarget*. First, in *NegativeIR* there exists only a single inflationary equilibrium following the central bank's intervention. This equilibrium coincides with the central bank's inflation target of 10%. Second, the key mechanism

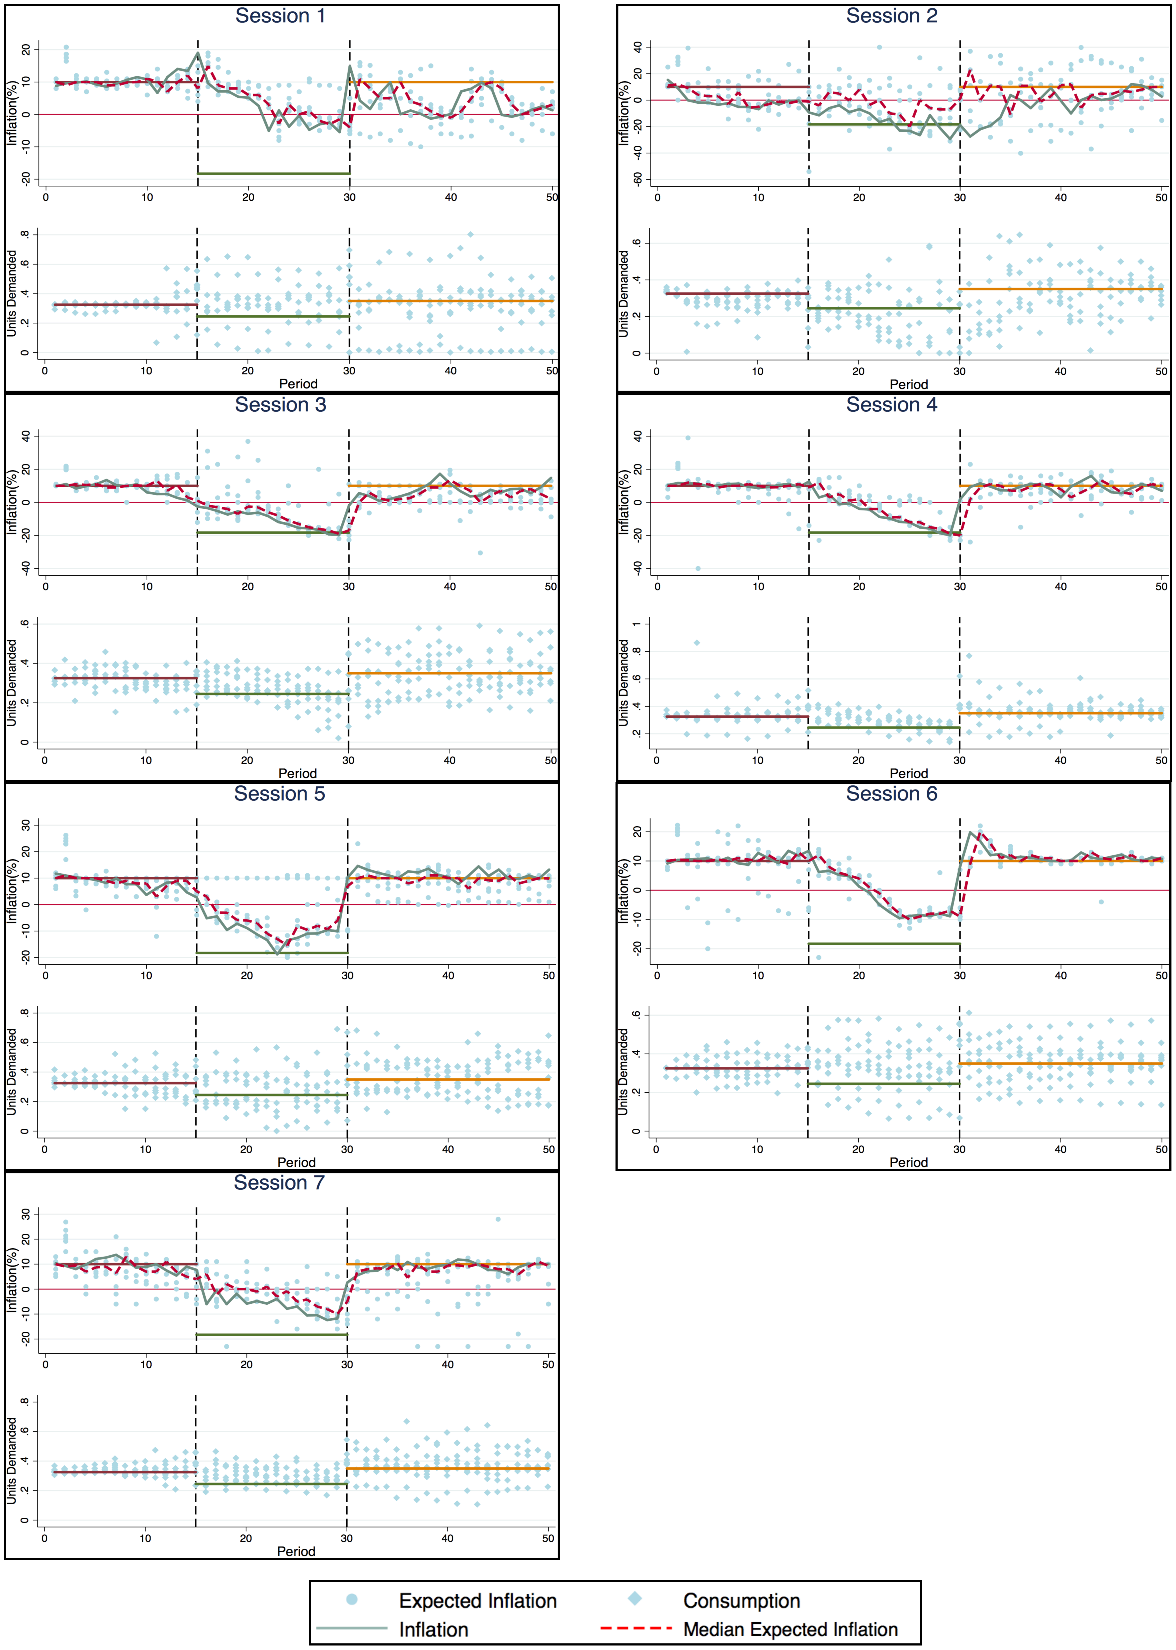

Figure 2.14: Aggregate inflation dynamics by session for *NegativeIR*

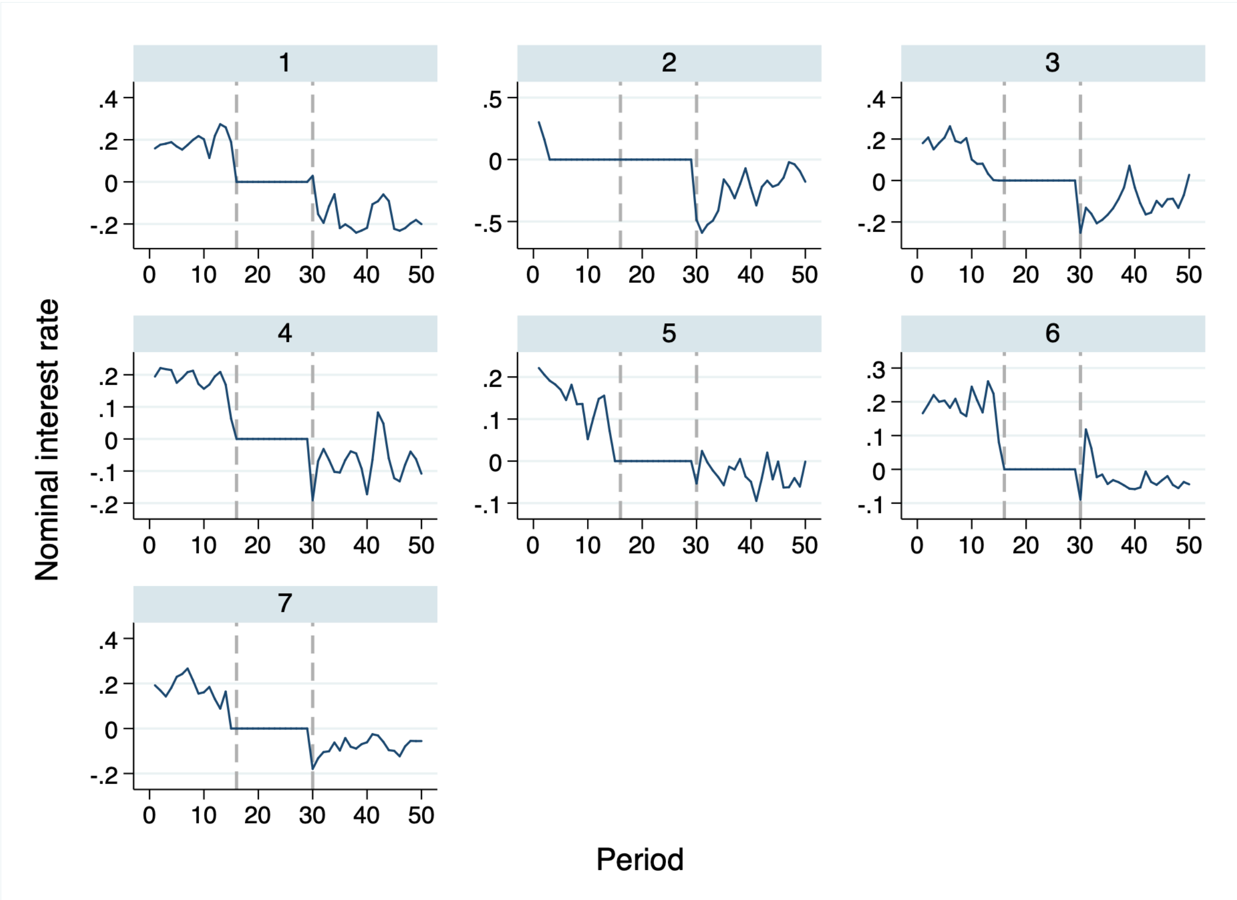

Figure 2.15: Nominal interest rates by session for *NegativeIR*

pushing the experimental economies out of secular stagnation in *NegativeIR* is a negative nominal interest rate that stimulates aggregate demand by increasing the appeal of current consumption relative to future consumption. That is, there is relatively less reliance on rational expectations to stimulate spending.

Eliminating the ZLB consistently stimulates aggregate spending and stabilizes our experimental economies at the targeted full-employment equilibrium. Though our experimental economy is devoid of some features that might reveal downsides to such an intervention (an ability for bank runs, for example) there is scope for agents to either misunderstand the implication of negative rates or, conversely, to experience a more severe wealth effect from negative interest rates and cut their spending further. We do not observe evidence of either

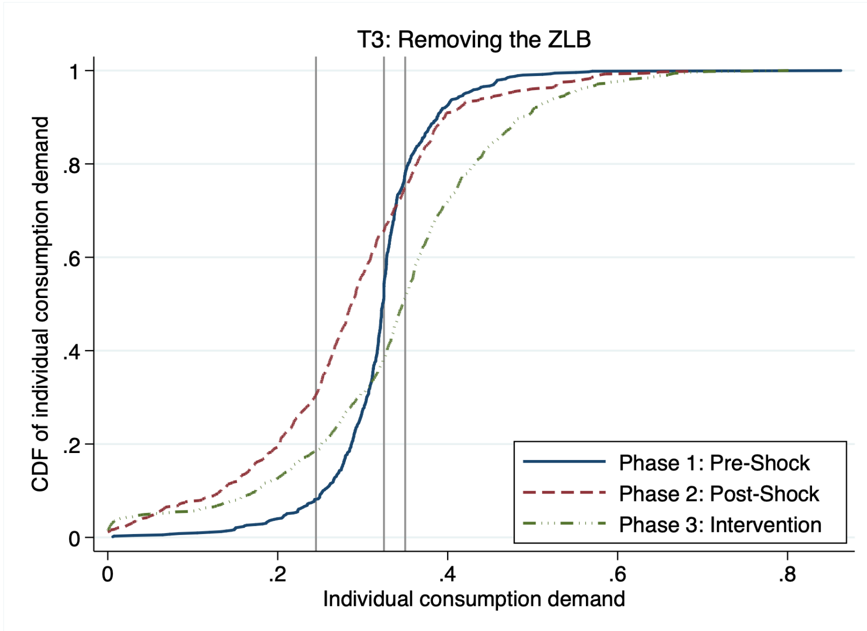

Figure 2.16: CDFs of consumption for all subjects and all periods for *NegativeIR*

Looking at Figure 2.15 gives a sense of the magnitude of negative rates used by the central bank to return these experimental economies to the full-employment equilibrium. We see in each experimental economy that nominal rates drop significantly in the intervention period and then revert almost instantly to a level much closer to zero. However, it is clear that, for most economies, the central bank is only capable of supporting the full-employment equilibrium by maintaining negative rates.

Figure 2.16 shows CDFs of consumption demand for *NegativeIR*. As was true with *Baseline* and *HigherTarget*, we again see that subjects consume optimally on average in Phase 1 and over consume in Phase 2. Now, unlike in *HigherTarget*, we see that consumption in Phase 3 is, on average, optimal. However, like in Phase 1 of all treatments, we see a large amount of heterogeneity in consumption.

We now evaluate the convergence of the economies, for each phase, to the predicted equilibria. In Phase 1, the predicted steady state inflation rate is 10%. We fail to reject H1 that inflation reached this level in *Baseline* and *NegativeIR* (Wilcoxon signed rank tests, $p = 0.26$ (N=8) in

Baseline, $p = 0.13$ (N=7) in *NegativeIR*.) In *HigherTarget*, inflation is slightly but significantly above the equilibrium at 13.2% ($p = 0.05$, N=9).

We hypothesized in H2 that large deleveraging shocks would inflation to stabilize at the secular stagnation equilibrium. We compare the mean inflation in the final five periods of Phase 2 to the equilibrium prediction of -24.4% inflation (in sessions that reached period 25). We reject H2 that the economies have converged to the secular stagnation equilibrium (Wilcoxon signed rank tests, $p = 0.03$ (N=6) in *Baseline* , $p = 0.008$ (N=9) in *HigherTarget*, and $p = 0.018$ (N=7) in *NegativeIR*.)

We also hypothesized that raising the inflation target to a sufficiently high level will move an economy to the targeted inflationary equilibrium of 30%. We compute the mean inflation in Phase 3 of *HigherTarget* after period 40. Mean inflation in Phase 3 after period 40 is -4.3%. We reject H3 that the economies converge to the new target inflationary equilibrium (Wilcoxon signed rank test, $p = 0.018$ (N=7)).

Finally, we hypothesized that eliminating the ZLB would return the economies in *NegativeIR* to their targeted inflationary equilibrium of 10%. Again, we compute the mean inflation in Phase 3 of *NegativeIR* after period 40. Across all sessions, mean inflation is 8.01%. We fail to reject H4 that the economies converge to the original inflationary equilibrium (Wicoxon signed rank test, $p = 0.31$ (N=7)).

Next, we can assess how the shocks and interventions alter the heterogeneity in subjects' forecasting heuristics. The different heuristic models we consider are listed below in Table 2.2. We focus our analysis on the forecasts related to concurrent inflation.

Model Class	Heuristic Name	Model
M1	Target	$E_{i,t}\pi_t = \pi_t^*$
M2	Naive Inflation	$E_{i,t}\pi_t = \pi_{t-1}$
M3	Constant Gain	$E_{i,t}\pi_t = E_{t-1}\pi_{t-1} + \gamma(E_{i,t-1}\pi_{t-1} - \pi_{t-1})$
M4	Trend-chasing	$E_{i,t}\pi_t = \pi_{t-1} + \tau(\pi_{t-1} - \pi_{t-2})$
M5	Naive Price	$E_{i,t}P_t = P_{t-1}$

Table 2.2: Inflation forecasting heuristics

M1 Target assumes that a subject bases her price forecast on the assumption that inflation today will equal the central bank's inflation target. Given the non-stochastic nature of the environment, this is a rational expectations equilibrium. M2 Naive Inflation assumes that a subject bases her price forecast on the assumption that inflation today will equal inflation yesterday. M3 Constant Gain assumes that a subject forms a price forecast today by updating yesterday's inflation expectation based on yesterday's inflation expectation error. Given this formulation, we consider a range of parameterizations of $\gamma \in [-1.5, -0.1]$. M4 Trend-chasing assumes that a subject's inflation forecast is an extrapolation of yesterday's inflation based on the recent trend in inflation. Given this formulation, we consider a range of parameterizations $\tau \in [0.1, 1.5]$. M5 Naive Price assumes that a subject forms a price expectation today based solely on yesterday's price.

We classify a subject by comparing, in each period, her price (or implied inflation) expectation for today to the price forecast arising from each of M1-M5. We then calculate the mean absolute error for each hypothetical heuristic (and for each parameter value for M3, M4) and classify participants as belonging to the heuristic that has the minimum RMSE.[15]

EMR's theory assumes that subjects are perfectly rational when forming price expectations. This, ex ante, seemed like an extreme assumption that could possibly form a wedge that would prevent the predictions of this model from mapping into reality. This assumption was relaxed by Gibbs (2017) who shows that the predictions of this are E-stable and thus survive, under a form of least squares learning.

We observe consistent heuristics in Phase 1 of all treatments. Trend-chasing is the dominant heuristic capturing 43-61% of participants' forecasting behavior. Forecasting according to the central bank's target and constant-gain heuristics describe a small minority of participants. We do not observe purely naive inflation forecasts, and only a small number of participants who forecast prices naively.

The deleveraging shock at the beginning of Phase 2 generates significant heterogeneity in heuristics, with all five classes of heuristics represented. Usage of the central bank's target de-

[15]Note that M3 is equivalent to M2 for $\gamma = -1$. In the case that participants were classified in both, we assign their type to be M2 Naive Inflation.

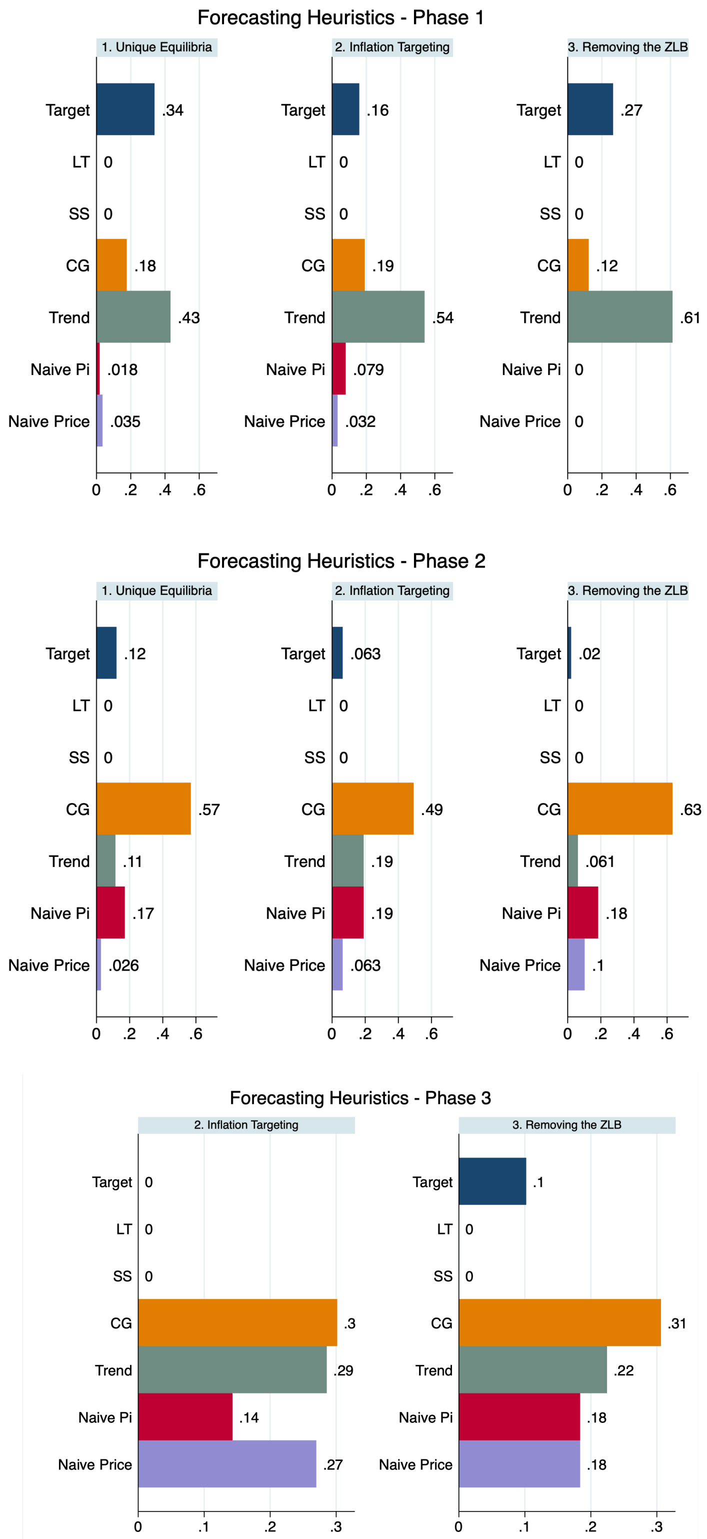

Figure 2.17: Forecasting heuristics by phase and by treatment

clines in all three treatments (from 34% to 12% in *Baseline*, 16% to 6.3% in *HigherTarget*, and 27% to 2% in *NegativeIR*), and is rational as the target is no longer an equilibrium outcome. Note that neither the liquidity trap or secular stagnation equilibria are focal points for participants. This is because we did not inform them of these equilibrium values. Nonetheless, participants' expectations do not adjust in line with these equilibria. Rather, constant-gain learning becomes the dominant heuristic as participants grapple with forecasting in an unfamiliar environment. This comes at a significant reduction in trend-chasing heuristics.

Increasing the inflation target in Phase 3 does not increase the share of participants using the central bank's target (the share falls to 0%). That is, no participant perceived the central bank's new inflation target of 30% as credible. The most striking change is the increase in share of participants simply forecasting based on last period's price, i.e. forecast zero percent inflation. We observe the proportion of naive price forecasts increase from 2.6% in Phase 2 to 27% in Phase 3, consistent with increased confusion about the environment. Trend-chasing heuristics also become more prevalent, with their share of the population rising from 11% in Phase 2 to 29% in Phase 3.

When the central bank eliminates the ZLB in *NegativeIR* more participants are willing to utilize the central bank's 10% inflation target as their forecast. This is still less than what we observe in Phase 1, suggesting that some of the credibility loss associated with the experience at the ZLB is permanent. The distribution of forecasting heuristics otherwise appears quite comparable to that of *HigherTarget*.

2.6 Conclusion

We experimentally investigate the ability of unconventional monetary policies to alleviate secular stagnation at the zero lower bound. Using the overlapping generations structure of Eggertsson, Mehrotra, Robbins (2019) as the model for our experimental economy, we find that EMR's model is robust to deviations from rationality along several dimensions. We find that the equilibria predicted by this model obtain even when economic agents form expectations using heuristics that are much less sophisticated than rational expectations and least-squares learning as explored by Gibbs (2017). High inflation equilibrium equilibria are relatively robust to deviations away from optimal consumption profiles. However, systematic over-consumption driven by Pigouvian wealth effects can prevent an economy reaching the secular stagnation equilibrium following a deleveraging shock.

We find that deleveraging shocks consistently produce results that qualitatively match the model's description and, across all treatments, often triggers secular stagnation. However, central bank interventions via inflation targeting consistently fails to pull economies out of deflationary traps. This result is consistent with Arifovic and Petersen (2017) who find that neither qualitative nor quantitative communication of higher inflation targets in a liquidity trap is sufficient to stimulate inflation expectations in a learning-to-forecast experimental environment. Removing the zero lower bound and allowing interest rates to become negative, on the other hand, reliably returns our experimental economies to the full-employment, inflationary equilibria. Our results suggest that policies aimed at stimulating aggregate demand through increased real wealth balances are more effective than those relying on rational expectations.

2.7 References

1. Ahrens, Steffen, Joep Lustenhouwer, and Michele Tettamanzi. "The stabilizing role of forward guidance: A macro experiment." (2017).

2. Arifovic, Jasmina. "Genetic algorithms and inflationary economies." Journal of Monetary Economics 36.1 (1995): 219-243.

3. Arifovic, Jasmina, and Luba Petersen. "Stabilizing expectations at the zero lower bound: experimental evidence." Journal of Economic Dynamics and Control 82 (2017): 21-43.

4. Arifovic, Jasmina, Stephanie Schmitt-Grohé, and Martín Uribe. "Learning to live in a liquidity trap." Journal of Economic Dynamics and Control 89 (2018): 120-136.

5. Ballinger, T. Parker, Michael G. Palumbo, and Nathaniel T. Wilcox. "Precautionary saving and social learning across generations: an experiment." The Economic Journal, 113.490 (2003): 920-947.

6. Benhabib, Jess, Stephanie Schmitt-Grohé, and Martin Uribe. “Monetary policy and multiple equilibria." American Economic Review 91.1 (2001): 167-186.

7. Bosch-Domènech, Antoni, and Joaquim Silvestre. “Credit constraints in general equilibrium: experimental results." The Economic Journal 107.444 (1997): 1445-1464.

8. Cadsby, Charles Bram, and Murray Frank. “Experimental tests of Ricardian equivalence." Economic Inquiry 29.4 (1991): 645-664.

9. Camerer, Colin. “Bounded rationality in individual decision making." Experimental Economics 1.2 (1998): 163-183.

10. Carbone, Enrica, and John Duffy. “Lifecycle consumption plans, social learning and external habits: Experimental evidence." Journal of Economic Behavior and Organization 106 (2014): 413-427.

11. Carbone, Enrica, and John D. Hey. “The effect of unemployment on consumption: an experimental analysis." The Economic Journal 114.497 (2004): 660-683.

12. Christiano, Lawrence, Martin S. Eichenbaum, and Benjamin K. Johannsen. “Does the New Keynesian Model Have a Uniqueness Problem?" No. w24612. National Bureau of Economic Research, 2018.

13. Cornand, Camille, and Paul Hubert. “On the external validity of experimental inflation forecasts: A comparison with five categories of field expectations." (2018).

14. Cornand, Camille, and Cheick Kader M’baye. “Does inflation targeting matter? An experimental investigation." Macroeconomic Dynamics (2016): 1-40.

15. Crockett, Sean, John Duffy, and Izhakian, Yehuda. “An experimental test of the Lucas asset pricing model." Review of Economic Studies, (forthcoming).

16. Del Negro, Marco, Marc P. Giannoni, and Frank Schorfheide. “Inflation in the great recession and new keynesian models." American Economic Journal: Macroeconomics 7.1 (2015): 168-96.

17. Di Maggio, M., Kermani, A., Keys, B. J., Piskorski, T., Ramcharan, R., Seru, A., & Yao, V. (2017). “Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging." American Economic Review 107.11 (2017): 3550-88.

18. Duffy, J. (2016), “Macroeconomics: A Survey of Laboratory Research,” in: J.H. Kagel and A.E. Roth (eds.), The Handbook of Experimental Economics Vol. 2, Princeton: Princeton University Press, pp. 1-90.

19. Eggertsson, Gauti B., and Paul Krugman. “Debt, deleveraging, and the liquidity trap: A Fisher-Minsky-Koo approach." The Quarterly Journal of Economics 127.3 (2012): 1469-1513.

20. Eggertsson, Gauti B., Neil R. Mehrotra, and Jacob A. Robbins. “A model of secular stagnation: Theory and quantitative evaluation." American Economic Journal: Macroeconomics, (forthcoming).

21. Eggertsson, Gauti B., Neil R. Mehrotra, and Lawrence H. Summers. “Secular stagnation in the open economy." American Economic Review 106.5 (2016): 503-07.

22. Eggertsson, G. B., Juelsrud, R. E., Summers, L. H., Wold, E. G. (2019). Negative nominal interest rates and the bank lending channel (No. w25416). National Bureau of Economic Research.

23. Fazzari, Steven M., James Morley, and Irina Panovska. “When Is Discretionary Fiscal Policy Effective?." (2018).

24. Evans, George, Guse, Eran, and Honkapohja, Seppo. “Liquidity traps, learning and stagnation." European Economic Review, 52.8, (2008): 1438-1463.

25. Fenig, Guidon, Mariya Mileva, and Luba Petersen. "Deflating asset price bubbles with leverage constraints and monetary policy." Journal of Economic Behavior Organization 155 (2018): 1-27.

26. Fenig, Guidon, and Luba Petersen. “Distributing scarce jobs and output: Experimental evidence on the dynamic effects of rationing." Experimental Economics 20.3 (2017): 707-735.

27. Filardo, Andrew, and Boris Hofmann. “Forward guidance at the zero lower bound." (2014).

28. Fréchette, Guillaume R. “Laboratory experiments: Professionals versus students." (2011).

29. Fréchette, Guillaume R. “Session-effects in the laboratory." Experimental Economics 15.3 (2012): 485-498.

30. Gabaix, Xavier. A Behavioral New Keynesian model. No. w22954. National Bureau of Economic Research, 2016. Gibbs, Christopher G. “Learning to believe in secular stagnation." Economics Letters 163 (2018): 50-54.

31. Gibbs, Christopher G. “Learning to believe in secular stagnation." Economics Letters 163 (2018): 50-54.

32. Hansen, Stephen, and Michael McMahon. “Shocking language: Understanding the macroeconomic effects of central bank communication." Journal of International Economics 99 (2016): S114-S133.

33. Haldane, Andrew G. “A Little More Conversation, A Little Less Action." Bank of England-Speech (2017).

34. Haldane, Andrew, and Michael McMahon. “Central Bank Communications and the General Public." AEA Papers and Proceedings. Vol. 108. 2018.

35. Hansen, Alvin H. “Economic progress and declining population growth." The American Economic Review 29.1 (1939): 1-15.

36. Heemeijer, P., Hommes, C., Sonnemans, J., Tuinstra, J. “Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation." Journal of Economic dynamics and control 33.5 (2009): 1052-1072.

37. Hey, John D., and Valentino Dardanoni. “Optimal consumption under uncertainty: An experimental investigation." The Economic Journal (1988): 105-116.

38. Hommes, Cars. “The heterogeneous expectations hypothesis: Some evidence from the lab." Journal of Economic Dynamics and Control 35.1 (2011): 1-24.

39. Hommes, Cars, Domenico Massaro, and Isabelle Salle. “Monetary and Fiscal Policy Design at the Zero Lower Bound: Evidence from the Lab." Workshop on Socio-Economic Complexity. Vol. 23. 2015.

40. Kessler, Judd, and Lise Vesterlund. “The external validity of laboratory experiments: The misleading emphasis on quantitative effects." Handbook of Experimental Economic Methodology, Oxford University Press, Oxford, UK. (2015)

41. Kryvtsov, Oleksiy, and Luba Petersen. “Expectations and monetary policy: experimental evidence". No. 2013-44. Bank of Canada Working Paper, 2013.

42. Kryvtsov, Oleksiy, and Luba Petersen. "Central Bank Communication that Works." Journal of Monetary Economics, forthcoming.

43. Lei, Vivian, and Charles N. Noussair. “An experimental test of an optimal growth model." American Economic Review 92.3 (2002): 549-570.

44. Lian, Peng, and Charles R. Plott. “General equilibrium, markets, macroeconomics and money in a laboratory experimental environment." Economic Theory 12.1 (1998): 21-75.

45. Lucas Jr, Robert E. “Asset prices in an exchange economy." Econometrica: Journal of the Econometric Society (1978): 1429-1445.

46. Lucas Jr, Robert E. “Adaptive behavior and economic theory." Journal of Business (1986): S401-S426.

47. Marimon, Ramon, and Shyam Sunder. “Indeterminacy of equilibria in a hyperinflationary world: experimental evidence." Econometrica: Journal of the Econometric Society (1993): 1073-1107.

48. Marimon, Ramon, and Shyam Sunder. “Expectations and learning under alternative monetary regimes: an experimental approach." Economic Theory 4.1 (1994): 131-162.

49. Marimon, Ramon, and Shyam Sunder. “Does a constant money growth rule help stabilize inflation?: experimental evidence." Carnegie-Rochester Conference Series on Public Policy. Vol. 43. North-Holland, 1995.

50. Meissner, Thomas. “Intertemporal consumption and debt aversion: an experimental study." Experimental Economics 19.2 (2016): 281-298.

51. Meissner, Thomas, and Davud Rostam-Afschar. “Learning Ricardian Equivalence." Journal of Economic Dynamics and Control 82 (2017): 273-288.

52. Mirdamadi, Michael, and Luba Petersen. “Macroeconomic Literacy and Expectations." (2018).

53. Mokhtarzadeh, Fatemeh, and Luba Petersen. “Coordinating expectations through central bank projections." (2017).

54. Nakov, A. “Optimal and Simple Monetary Policy Rules with Zero Floor on the Nominal Interest Rate." International Journal of Central Banking (2018).

55. Noussair, Charles N., Damjan Pfajfar, and Janos Zsiros. “Pricing decisions in an experimental dynamic stochastic general equilibrium economy." Journal of Economic Behavior Organization 109 (2015): 188-202.

56. Noussair, Charles, and Kenneth Matheny. “An experimental study of decisions in dynamic optimization problems." Economic Theory 15.2 (2000): 389-419.

57. Noussair, Charles, Charles Plott and Raymond Riezman. “An Experimental Investigation of the Patterns of International Trade”. American Economic Review, June 1995, pages 462-491.

58. Noussair, Charles, Charles Plott and Raymond Riezman. “The Principles of Exchange Rate Determination in an International Finance Experiment”, Journal of Political Economy, (1997), pages 822-862.

59. Noussair, Charles, Charles Plott and Raymond Riezman. “Production, Trade and Exchange Rates in Large Experimental Economies”, European Economic Review, (2007), Vol. 51, No. 1, pages 49-76.

60. Ochs, Jack. “Coordination problems." Handbook of experimental economics (1995): 195-252.

61. Offerman, Theo, Jan Potters, and Harrie AA Verbon. “Cooperation in an overlapping generations experiment." Games and Economic Behavior 36.2 (2001): 264-275.

62. Patinkin, Don. "Price flexibility and full employment." American Economic Review 38.4 (1948): 543-564.

63. Samuelson, Paul A. “An exact consumption-loan model of interest with or without the social contrivance of money." Journal of Political Economy 66.6 (1958): 467-482.

64. Petersen, Luba. “Do expectations and decisions respond to monetary policy?" Journal of Economic Studies 42.6 (2015): 972-1004.

65. Summers, Lawrence H. “US economic prospects: Secular stagnation, hysteresis, and the zero lower bound." Business Economics 49.2 (2014): 65-73.

66. Schmitt-Grohé, Stephanie, and Martín Uribe. “Liquidity traps and jobless recoveries." American Economic Journal: Macroeconomics 9.1 (2017): 165-204.

67. Van der Heijden, E. C., Nelissen, J. H., Potters, J. J., Verbon, H. A. (1998). “Transfers and the effect of monitoring in an overlapping-generations experiment." European Economic Review, 42(7), 1363-1391.

68. Woodford, Michael. “Learning to believe in sunspots." Econometrica: Journal of the Econometric Society (1990): 277-307.

69. Wu, Jing Cynthia, and Fan Dora Xia. “Measuring the macroeconomic impact of monetary policy at the zero lower bound." Journal of Money, Credit and Banking, 48.2-3 (2016): 253-291.

3. CENTRAL BANK DENSITY FORECASTS: DO HIGHER-ORDER MOMENTS MATTER?

Central banking has undergone a decades-long transparency revolution that brought about an increased reliance on communication to manage the expectations of economic agents to achieve policy objectives (Blinder et al., 2008). Communication is a particularly important tool whenever policy rates are constrained at the effective lower bound, which has been the case in the past decade across the world. Arguably, the first link in the context of monetary policy transmission is the effect of the communications on the market interest rates which are relevant for consumption and investment decisions in the economy — a topic we study in this paper. More specifically, we characterize the evolution of economic outlook across the monetary policy cycle and quantify the effects on the market interest rates and survey forecasts.

A primary component of a central bank's communication strategy is its projections of the key macroeconomic variables, including inflation and output. These summarize the conditions to which a central bank is reacting through its policy instruments. Many central banks publish density forecasts alongside their point predictions, conveying the central bank's uncertainty and outlook on the balance of risks. The Bank of England (BoE) was the first to publish 'fan charts' of its macroeconomic projections in 1998, which is now standard practice across the board.

There is a wide literature on monetary policy communication, mainly focusing on the publication of point forecasts. The main conclusion of this literature is that communication facilitates policy and acts as a coordination device for expectation formation. Swanson (2006), Blinder et al. (2008), Hubert (2014, 2015) support this finding with empirical analyses, while Kryvtsov and Petersen (2020) and Ahrens et al. (2019) provide experimental support. There is further evidence that central bank communication affects the market interest rates (see Gurkaynak, Sack, and Swanson, 2004 and Andrade and Ferroni, 2020 among others), making it relevant for policymaking. The communication of economic outlook has also proven to be important under the information channel of monetary policy discussed in seminal works of Nakamura and Steinsson (2018) and as well as Miranda-Agrippino and Ricco (2018), Hubert (2019), Hoesch et al. (2020), among others.

Though the literature on central bank communication is rather extensive, the use of density forecasts and the information they contain has been rather scant. By observing the various central bank behaviors, it also seems that there is no consensus in policy circles on what the communication of uncertainty delivers. As Petersen and Rholes (2020) point out, with the COVID-19 crises the central banks around the world have moved to different directions — for example, both the Federal Reserve and the BoE abandoned predictive densities in exchange for scenario analysis at the onset of the pandemic. On the other hand, in the April 2020 Monetary Policy Report, the Bank of Canada rid itself of point predictions and instead published only ranges. Some central banks, such as the European Central Bank, made no change to their communication strategy. The second interesting observation is that communication of uncertainty appears to be deemed important since, over time, the central banks are moving to a more timely release of that information. For instance, per Reuters article on November 5, 2020 "The Federal Reserve will publish new color around policymaker outlooks for interest rates and the economy, and release some details earlier, changes that could give fresh insight into rate-setting decisions ... [the changes should] provide a timely perspective on the risks or uncertainties that surround the modal or baseline projections," Powell said, "thereby highlighting some of the risk management considerations that are relevant for monetary policy." The heterogeneity in nature of central bank projections throughout the current global pandemic highlights the lack of consensus about when, how, or why to use predictive densities. Thus, it is not obvious how or even if information regarding risk and uncertainty contained in density forecasts matters from a policy perspective, even though density forecasts could conceivably influence financial markets and economic activity.

Despite this widespread publication of density forecasts, there is very limited evidence on how higher-order moments matter from a policy perspective. Rholes and Petersen (2020) show experimentally that communicating uncertainty alongside a point projection of inflation can increase individual-level forecast errors and forecast uncertainty relative to communicating only point projections. Hubert and Maule (2020) show empirically that private expectations respond strongly to the BoE's signals about future economic activity as conveyed in its Quarterly Inflation Report (QIR), suggesting that density forecasts might operate through a signaling channel. Hansen, McMahon, and Tong (2019) uses textual analysis of the BoE's QIR to show that the economic uncertainty conveyed by text (and which is orthogonal to information conveyed numerically at the 2-year forecast horizon) can have increasingly large effects along the yield curve. Further, the heterogeneity in the nature of central bank projections throughout the current global pandemic highlights the lack of consensus about when, how, or why to use predictive densities. Thus, it is not obvious how or even if information regarding risk and uncertainty contained in density forecasts matters from a policy perspective, even though density forecasts could conceivably influence financial markets and economic activity.

The goal of this paper is to characterize and understand whether and how the financial markets and private sector expectations respond to the revisions of higher-order moments of the economic outlook. We use the data from the BoE because the BoE provides the longest time-series of relevant data for an empirical study. We rely on a high-frequency identification approach, taking the revisions from one publication to the next, announced on a specific calendar day to identify the effects of central bank economic outlook (revisions) on the UK term structure of interest rates and the expectations of professional forecasters. The paper closest to our work is Hansen et al. (2019) who use textual analysis as well as density forecast updates to understand the long interest rate sensitivity to macroeconomic news (relative to the short end of the yield curve). The main difference between their work and ours is that we differentiate between uncertainty and skewness to better understand the various properties that are typically important in considering density forecasts.

We find that revisions to higher-order moments matter more than revisions to first-order moments for interest rate dynamics. In fact, there is not much action due to the first moment movements. On the other hand, an increase in the higher moments of output growth densities inverts the yield curve, typically known to forecast a recession. On the other hand,an increase in inflation uncertainty does not seem to affect the financial markets much, while skewness appears to be important. Moreover, the effect of higher moments is state-dependent — uncertainty revisions play through in expansions, while skewness revisions provide informative signals in contractions. Historical decomposition suggests yields have responded to higher-order moments for many decades, and the response was strong during the financial crisis.

When looking at the Blue Chip Financial Forecasts, we find that the three-, six- and twelve-month ahead consensus forecasts (and forecast disagreement) of short- and long-term interest rates are strongly correlated with revisions to uncertainty, while revisions to the balance of risks, expressed through revisions to skewness, are not.

Our paper is organized as follows. Section 2 presents our key data sets and our data transformations and introduces the institutional details for the BoE communication through the monetary policy cycle. Section 3 lays out our identifying assumptions and estimation strategy. Section 4 presents our findings while Section 5 concludes.

3.1 Data

The BoE's nine-person Monetary Policy Committee (MPC) began publishing density forecasts, termed 'fan charts', of inflation, output growth, and unemployment in its QIR in 1997.[1] The MPC, which adheres to an inflation-targeting regime adopted in the U.K. 1993, publishes fan charts to convey the inherent uncertainty surrounding its economic outlook and to provide its collective outlook on the balance of risks. These fan charts present deciles of subjective estimates of the probability distribution of the Bank's forecast of each of these key macroeconomic variables (Eler, 2005; Mitchell and Weale, 2019).[2] We provide an example of these projections for GDP growth, inflation, and unemployment in Figure 3.1.

The MPC constructs its fan charts using a split-normal distribution with a common mode but two different variances.[3] The BoE' particular representation centers the distribution at the mode (μ) with an uncertainty (σ^2), while the skewness (ξ) controls the relative behavior of the two halves of the distribution. Thus, the probability density function (pdf), $f(x)$, is given by

[1]The BoE, in conjunction with the National Institute of Economic and Social Research, began publishing a measure of uncertainty surrounding inflation forecasts as early as February 1996.

[2]The QIR itself includes graphical depictions of these fan charts. However, the BoE makes available the numerical information used to construct these fan charts at www.bankofengland.co.uk.

[3]These values are identical whenever the distribution is symmetric about the mode, which will equal the mean.

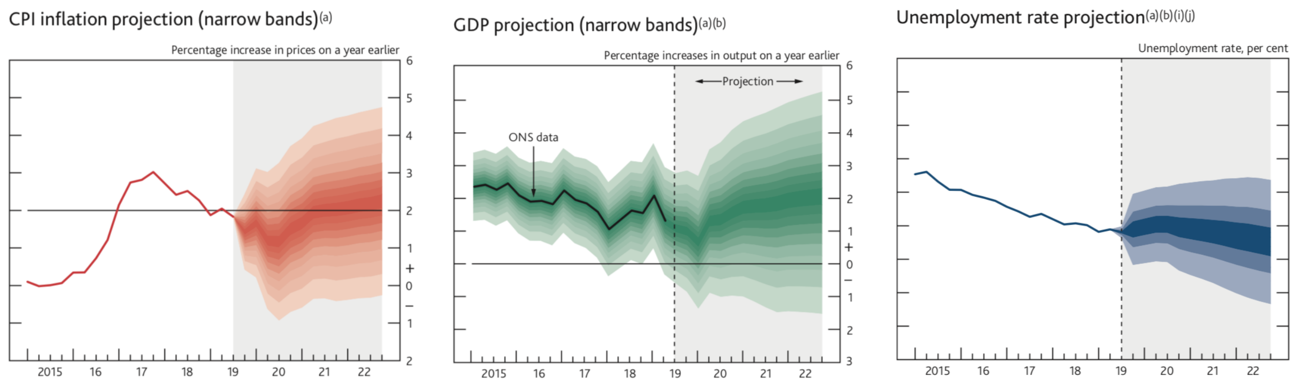

(a) Inflation (b) Output Growth (c) Unemployment Rate

Figure 3.1: Density forecasts from the Bank of England's Monetary Policy Report.

$$f(x) = \begin{cases} Ae^{\left[\frac{-(x-\mu)^2}{2\sigma_1^2}\right]}, x \leq \mu \\ Ae^{\left[\frac{-(x-\mu)^2}{2\sigma_2^2}\right]}, x \geq \mu \end{cases} \tag{3.1}$$

where the following definitions hold

- $\sigma^2 = (1+\gamma)\sigma_1^2 = (1-\gamma)\sigma_2^2,$
- $E(X) = \mu + \sqrt{\frac{2}{\pi}}(\sigma_2 - \sigma_1)$
- $var(X) = (1 - \frac{2}{\pi})(\sigma_2 - \sigma_1)^2 + \sigma_1\sigma_2$
- $A = \left(\frac{\sqrt{2\pi}(\sigma_1+\sigma_2)}{2}\right)^{-1}$
- $\gamma^2 = 1 - 4\left(\frac{\sqrt{1+\pi\xi^2/\sigma^2}-1}{\pi\xi^2/\sigma^2}\right)^2$

and parameter values are specific to forecast horizons and forecast date. That is, the MPC assigns potentially unique values of $\sigma,\ \gamma$ for each forecasting horizon at each time t. For example, when forecasting period $t+k$ while in period t, the MPC sets values of $\sigma_{t,t+k},\ \gamma_{t,t+k}$ for $k \in K \equiv \{1, 2, ..., K\}$, where $\sigma_{t,t+k}, \sigma_{t,t+j},\ j, k \in K$ can differ and $\sigma_{t,t+k}, \sigma_{(t-l),(t+k+l)}$ can differ.

Our focus in this paper is on how both the UK's yield curve and private expectations respond to revisions to the BoE's outlook on uncertainty (σ, the second central moment) and risk (γ, the third central moment).[4] For comparative purposes, we also consider how yields and expectations respond to revisions of the point forecasts (μ, the first central moment).

An intuitive way to think of σ is as the uncertainty surrounding the MPCs point projections of some key economic variable. Visually, an increase in this parameter creates a more diffuse set of uncertainty bands surrounding the MPC's point projection[5] One can think of γ as a measure of the MPC's subjective outlook on the balance of risks associated with some key economic variable. Whenever $\gamma = 0$, the split-normal density function reduces to the pdf of the normal distribution. Positive values of γ indicate that upside risk dominates downside risk. Visually, this would mean the MPC's density forecast is positively-skewed. The opposite is true for negative values of γ.[6]

To get a sense of how the MPC sets $\sigma,\ \gamma$, and how its approach might have changed over time, we consider the MPC's outlook on uncertainty and risk surrounding both the 2008 collapse of Lehman Brothers and the 2016 'Brexit' referendum.

When Lehman Brothers filed for bankruptcy on September 15, 2008, the firm held more than \$700 billion in liabilities, which obviates why the failure of Lehman's shocked global financial markets. The MPC addressed Lehman's collapse directly in its QIR in November 2008: "A number of institutional failures, and in particular, the collapse of Lehman Brothers on 15 September, led to rising anxieties about the survival of other financial institutions internationally." This sense of 'rising anxieties' is reflected by the MPC's upward revision of GDP growth and inflation uncertainty measures at all forecast horizons in its 2008Q4 inflation report, which we show in Figure 3.3. The MPC responded to this event by revising its balance of risks so that its density forecasts of GDP growth and inflation were exactly symmetric.

Less than a decade later, on June 23, 2016, the United Kingdom's electorate voted to leave the European Union. The MPC expressed concern that the vote to leave the EU generated considerable economic uncertainty and argued in the August 2016 QIR that short- and medium-term output growth would fall as a consequence. Additionally, the MPC focused on the devaluation of the

[4] γ is the difference between the mean μ and the mode of the distribution.

[5] Note that σ is equal to the standard deviation whenever the distribution is symmetric (i.e. $\gamma = 0$)

[6] Wallis (2004) provides a more in-depth discussion of the BoE's density forecast.

pound sterling and the prospect of ensuing supply constraints. These concerns prompted the MPC to increase its level of forecast uncertainty about GDP growth, relative to its May 2016 report. Regarding inflation, the MPC made clear that the declining exchange rate was a secondary concern relative to impending supply constraints and dampened growth projections. The MPC, charged with maintaining full output and price stability, faced a trade-off between an inflationary policy and recessionary pressures. Ultimately, the MPC declared a willingness to endure temporary inflation above its 2% target to boost demand and supply to avoid larger declines in output growth. Thus, the MPC decided to cut the policy rate by 25 basis points following the electorates vote to leave the EU. These considerations and the MPC's policy action is reflected by the MPC's decision to leave unchanged its GDP growth skewness parameters while increasing skewness parameters for its short- and medium-term inflation forecast.

3.1.1 Timing

The MPC sets policy monthly and conveys its economic outlook quarterly (in February, May, August, and November) in its QIR. From November 1997 through May 2015, the BoE published the QIR one week after announcing information about that month's monetary policy decision. Since August of 2015, the BoE has released information about monetary policy decisions and its economic outlook simultaneously in the QIR.[7] Thus, we can consider the impact of uncertainty and skewness shocks independent of information about the BOE's policy decisions through the second quarter of 2015 but not after. Figure 3.2 provides an overview of the BoE's information release schedule.

3.1.2 Density forecast data

Density forecast data for inflation from 2004 and for GDP growth from 2007 are publicly available as part of the BoE's quarterly inflation report.[8] Density forecast data for inflation from 1997 through 2005 and for GDP growth from 1997 through 2007 are available through the U.K.'s

[7]The BoE began referring to this joint release as its Monetary Policy Report in November 2019.

[8]This data is available for download here: https://www.bankofengland.co.uk/inflation-report/inflation-reports

	Current schedule	New schedule
Wednesday		"Pre-MPC" meeting with staff presentations to the MPC (joint MPC-FPC briefing meetings will take place four times a year).
Thursday		Stage 1: MPC deliberation meeting
Friday	"Pre-MPC" meeting with staff presentations to the MPC	
Monday		Stage 2: MPC policy discussion meeting
Tuesday		
Wednesday	Stage 1: MPC deliberation meeting	Stage 3: MPC decision meeting
Thursday	• Stage 2: MPC discussion and decision meeting • Announcement of monetary policy decision	• Announcement of monetary policy decision and simultaneous publication of minutes • Inflation Report and press conference in Inflation Report months
Wednesday one week later	Inflation Report and press conference in Inflation Report months	
Wednesday two weeks later	Minutes published	

Figure 3.2: Policy and Communication Schedule

national archive.[9]. Density forecasts of inflation from 2004 and onward use the consumer price index (CPI) to measure inflation while older inflation data uses retail purchases excluding mortgage payments (RPIX) to measure inflation.

The BoE publishes two different density forecasts for both GDP growth and inflation, each following a different assumption about the interest rate path. The bank forms one set of density forecasts by assuming that the prevailing nominal interest rate will continue throughout the forecast horizon. We call this the constant-rate assumption. The second set of density forecasts assumes the nominal interest rate matches the market's expected nominal rate throughout the forecast hori-

[9]This data is available for download here: https://webarchive.nationalarchives.gov.uk/20170831105150/ http://www.bankofengland.co.uk/publications/Pages/inflationreport/irprobab.aspx

zon. We call this the market-rate assumption. The BoE formed projections under the constant-rate assumption for 9 quarters, including a current-quarter forecast, from the beginning of our sample through May 2013 and for 13 quarters thereafter. The bank makes projections under market assumptions, when available, out to 13 quarters.

These two alternative interest rate assumptions typically yield different values of μ (point forecast) but have no impact on σ (uncertainty) or γ (skewness). Thus, we need not discern between these two interest rate assumptions when considering how rates and expectations respond to higher-order forecast moments.

We show the first moment of the BoE's density forecasts of GDP growth and inflation formed using both interest rate assumption in Figure 3.3, and σ and γ for both GDP growth and inflation in Figure 3.4.

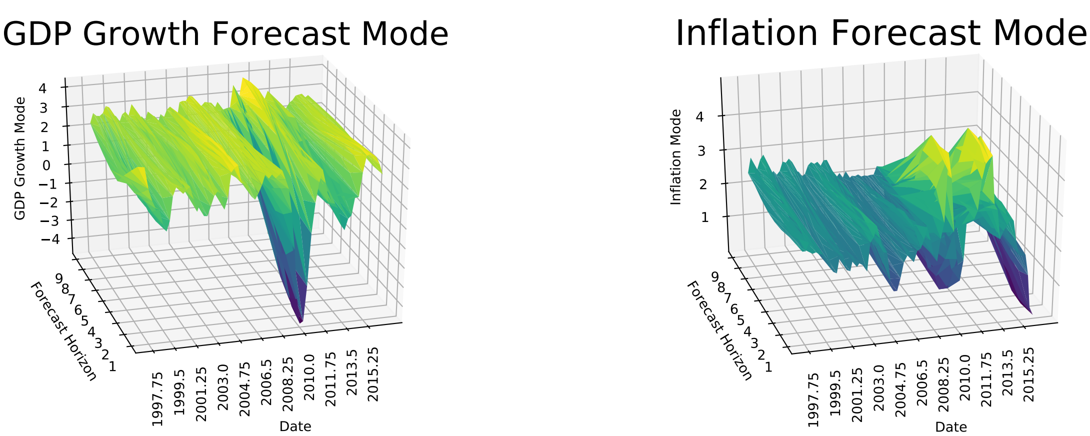

(a) (b)

Figure 3.3: First moments of the BoE's density forecasts of GDP growth and inflation

Note that movements in σ and γ are highly correlated over time across their respective forecast horizons. Also worth noting is that σ_π, σ_Y, for all forecast horizons, both exhibit considerably more variability before and through the early parts of the Great Recession than in subsequent forecasts, suggesting that MPC may have changed its approach to parameter selection. This change seems to correspond to the peak of a drastic increase in uncertainty surrounding density forecasts of both GDP growth and inflation that occurred following the global financial crisis.

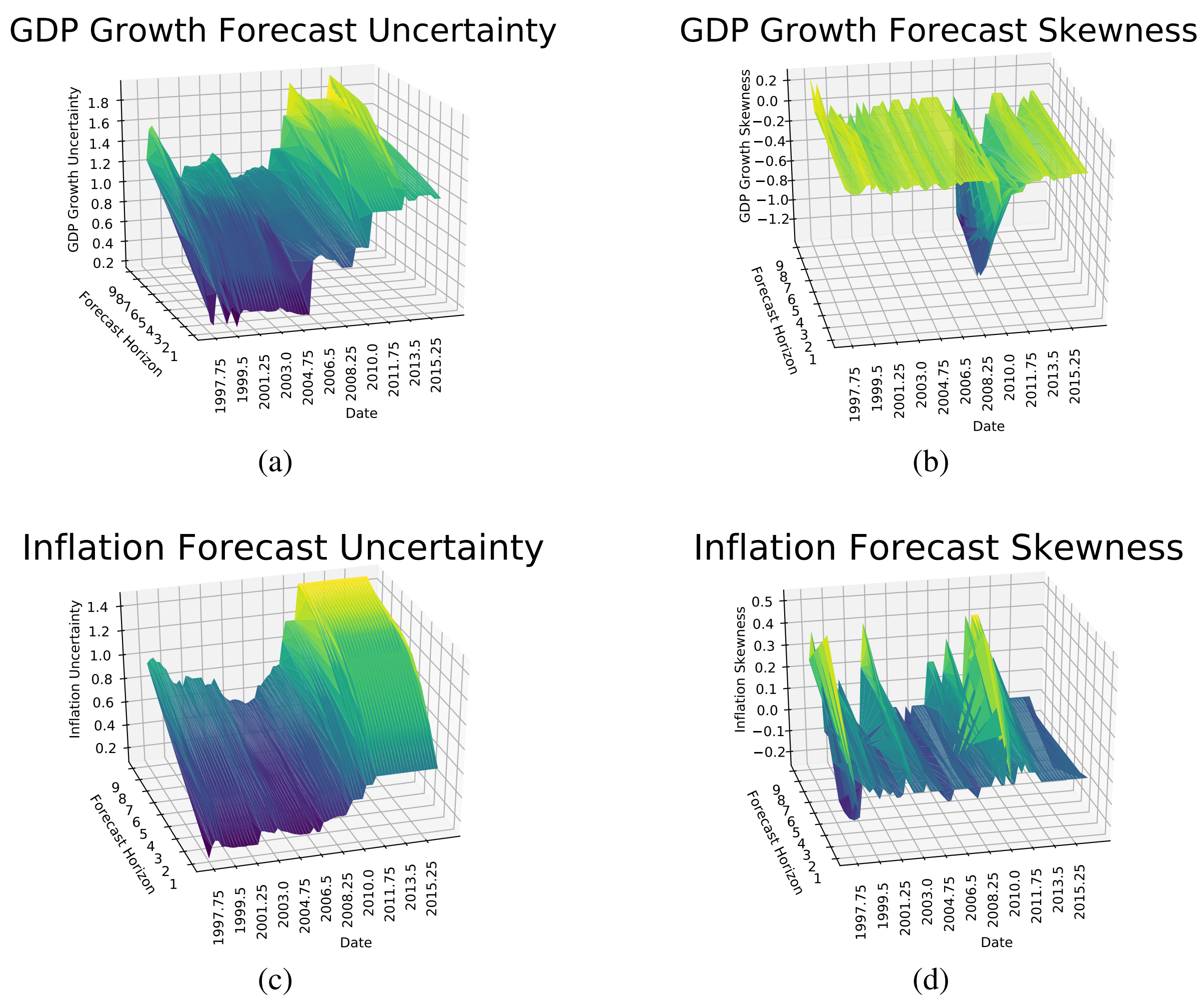

Figure 3.4: Uncertainty and skewness of GDP growth and inflation

Notice also that γ_π rarely takes on positive values at any forecast horizon. This means the BoE rarely expects downside risk to dominate across our time sample. Also interesting to note is that γ_Y doesn't behave similarly across the two most recent recessions, which aligns with the idea that the MPC has begun taking a much different approach to how it selects values of γ and σ when forming its density forecasts. Finally, it is also noteworthy that neither of γ_π nor γ_Y appears to be strongly counter-cyclical, which is true of most other measures of economic uncertainty (Bloom, 2014).

3.1.3 Daily data for the UK's yield curve

The BoE forms two different daily estimates of maturities along the UK's yield curve. The first estimate is for maturities ranging from one to 60 sixty months in one-month intervals. Following the BoE's nomenclature, we refer to this as the 'short-end' estimate. The second estimate is for maturities ranging from six months to 25 years in six-month intervals. The BoE has more recently begun estimating maturities out to 40 years. However, we only consider yields out to 25 years since these estimates are available for our entire time sample. We refer to this as the 'long-run' estimates. Data for both estimates are publicly available through the BoE's website.[10]

The BoE bases its daily estimates of maturities on UK government bonds (gilts) and yields in the general collateral(GC) repurchase agreement (repo) market. The bank generates synthetic zero-coupon bonds from GC repo rates to improve its estimates of shorter maturities, where gilts tend to be less liquid. The BoE uses the variable roughness penalty (VRP) method, based on the spline-based technique proposed by Waggoner (1997), to estimate the yield curve.

3.1.4 Data Transformations

Revisions

We obtain our main results using a high-frequency identification strategy. We project changes in maturities across the UK's yield curve that occur in a small window surrounding QIR releases onto shocks to the first three central moments of the BoE's density forecasts of GDP growth and inflation. Thus, we transform our data to obtain two primary components: a left-hand-side variable capturing these maturity changes and right-hand-side variables capturing shocks to our moments of interest for inflation and GDP growth (i.e. information shocks).

To obtain our information shocks, we exploit the fact that the BoE often revises the value of our parameters of interest used to forecast economic values for some fixed point in time in sequential

[10] This data is available for download at https://www.bankofengland.co.uk/statistics/yield-curves

	2010Q1	2010Q2	2010Q3	2010Q4	2011Q1	2011Q2	2011Q3	2011Q4	2012Q1	2012Q2
2010Q1	H0	H1	H2	H3	H4	H5	H6	H7	H8	
2010Q2		Z0	Z1	Z2	Z3	Z4	Z5	Z6	Z7	Z8

Figure 3.5: Sequential forecast example

Table of Correlations for Revision Factors						
	Y_μ	Y_{σ^2}	Y_ξ	π_μ	π_{σ^2}	π_ξ
Y_μ	1.00					
Y_{σ^2}	-.3170	1.00				
Y_ξ	.0088	.1773	1.00			
π_μ	-.0038	-.1795	.1413	1.00		
π_{σ^2}	-.0885	.1552	-.0399	.0727	1.00	
π_ξ	-.3030	.2649	.1937	.0311	-.0354	1.00
Prop.	.6	.67	.86	.69	.66	.83

Table 3.1: Table of correlations for information shocks

forecasts. For example, consider the BoE's forecasts of some key economic variable made in the first and second quarters of 2010, which we depict in Section 3.1.4.

Given these two forecasts, we first compute the change in parameter values $\mu,\ \sigma,\ \gamma$ associated with the BoE's density forecasts of inflation and output growth for each possible forecast horizon. Using values depicted in Section 3.1.4, we compute these revisions for the second quarter of 2010 as $\Delta Z_0 = Z_0 - H_1$, ΔZ_1, ΔZ_2, ..., ΔZ_7. Next, we use principle component analysis (PCA) to extract the first principle component from ΔZ_0, ΔZ_1, ΔZ_2, ..., ΔZ_7 to obtain $PC_{x,j,t}$ where $x = \{\mu, \sigma^2, \xi\}, j = \{\pi, Y\}$. This approach allows us to summarize the majority of the variation contained in our information shocks while abstracting away from certain horizon-specific idiosyncrasies. We graph these information shocks for GDP growth and inflation in Figure 3.6.[11] Section 3.1.4 gives the level of correlation between these factors.

We follow Diebold and Li (2006) to extract time-varying level, slope, and curvature factors from our daily yield estimate that models the UK's yield curve. Because we have daily data, we

[11] We provide similar graphs for levels in **??** our appendix.

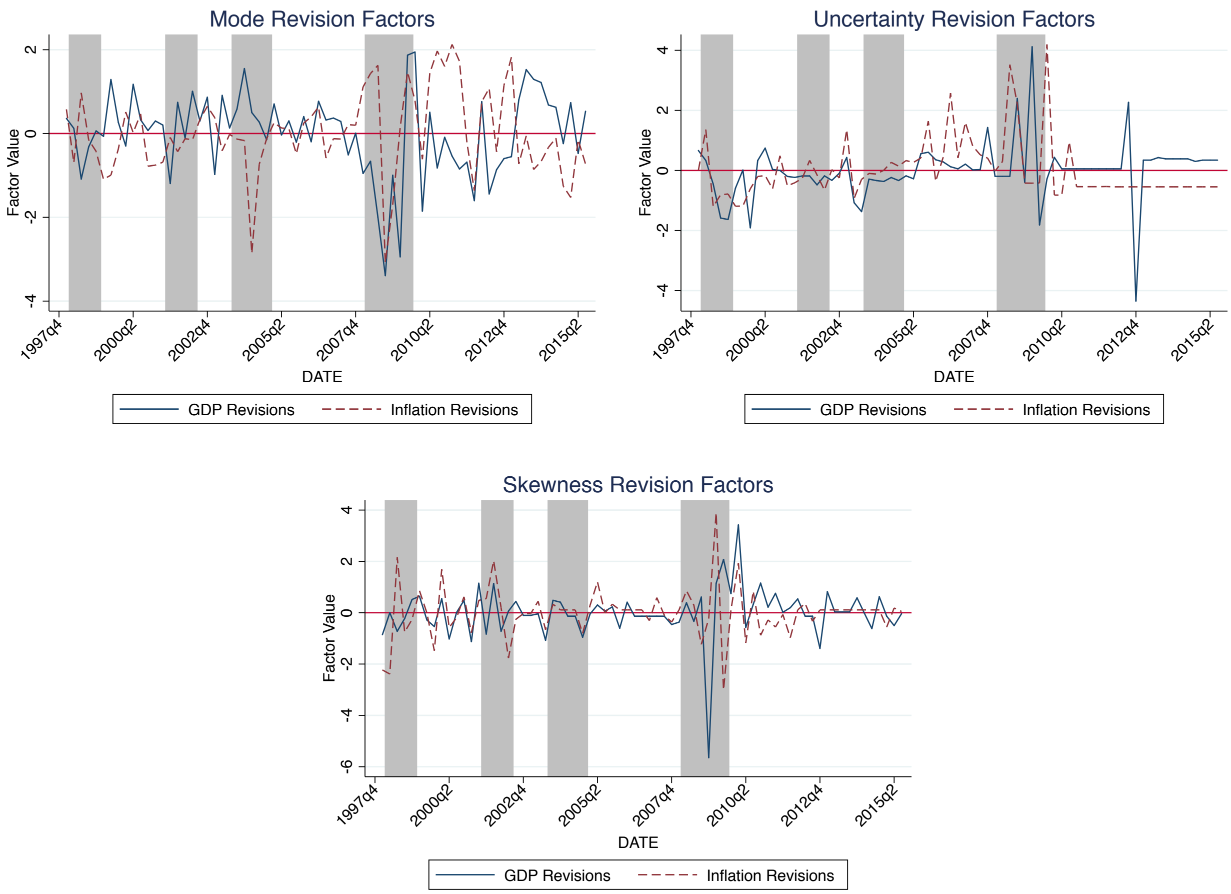

Figure 3.6: GDP growth and inflation revision factors

obtain daily estimates of these factors and can use the change in these estimated factors that occurs on inflation report release dates as outcome variables in the high-frequency identification scheme described above. To extract these factors, we use all available yield curve data ranging from the last quarter of 1997 through the second quarter of 2015. For maturities ranging from one to 60 months, we use data from the BoE's short-end yield curve estimates. These estimates are available for available for maturities at monthly intervals. For maturities further into the term structure, we use the BoE's estimates of the full yield curve. These estimates are available for maturities ranging from six months to 25 years in six-month intervals. Using this data, we estimate Equation (3.2)

$$yield_t(m) = \beta_{1,t} X_1 + \beta_{2,t} \left(\frac{1 - e^{\lambda_t m}}{\lambda_t m} \right) + \beta_{3,t} \left(\frac{1 - e^{\lambda_t m}}{\lambda_t m} - e^{\lambda_t m} \right). \quad (3.2)$$

Interpretations of each $\beta_{i,t}$ that follows directly from the behavior of the corresponding loadings $X_{i,m}$. We show both the factors and the loadings in Figure 3.7.

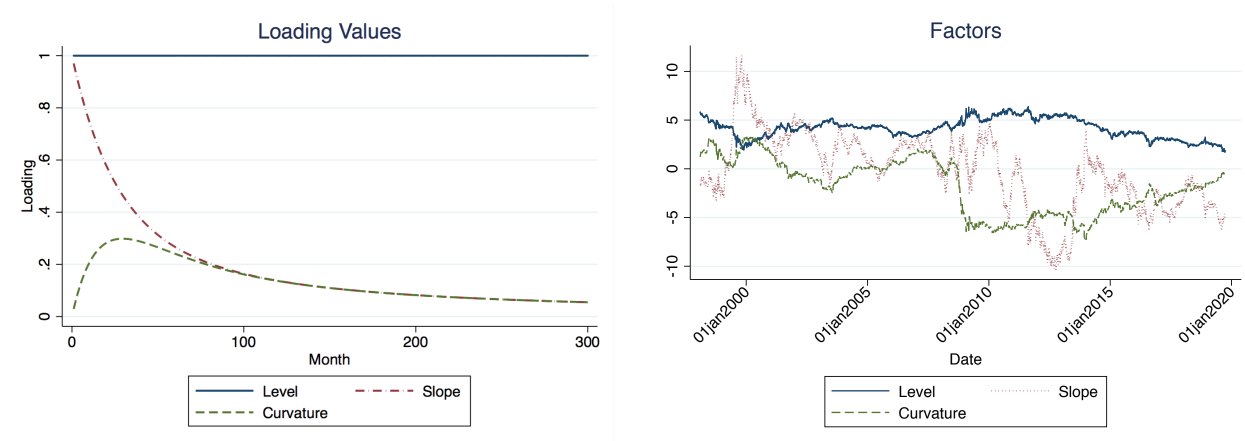

(a) Factor Loadings (b) Factors

Figure 3.7: Yield curve factor loadings.

The loading on $\beta_{1,t}$, X_1, is a constant so that we can interpret $\beta_{1,t}$ as our long-term (i.e. level) factor. $X_{2,m}$ decays monotonically from one to zero and so we can interpret $\beta_{2,t}$ as our short-term (i.e. slope) factor. $X_{3,m}$ begins at zero, peaks at around 30 months, and then converges monotonically toward zero. Thus, we treat $_{3,t}$ as the medium-term (i.e. curvature) factor.

3.2 Results

This section presents our main results, wherein we use a high-frequency identification approach, developed by Cook and Hahn (1989), Kuttner (2001), and Cochrane and Piazzesi (2002). For examples of this identification approach, see Nakamura and Steinsson (2018) or Gurkanayak, Sack, and Swanson (2004).

In our context, this identification scheme involves projecting changes in yield curve factors that occur in the 24-hour window surrounding QIR releases onto our information. Our identifying assumption is that variation in our outcomes of interest that occur in our tight window surrounding the release of the inflation report is driven by the information shocks we extract from sequential QIR releases.

Using $\Delta\beta_{i,t}$ and $\Delta PC_{x,j,t}$, where $i = \{l, c, s\}$ we estimate:

$$\Delta\beta_{i,t} = \alpha_{i,t} + \sum_{x,j} \psi_{x,j} PC_{x,j,t} + \kappa_{i,t} FTSE_{t-1} + \epsilon_{i,t} \tag{3.3}$$

which we estimate using a Newey-West estimator. Here, $\psi_{x,j}$ captures the causal relationship between our information shocks and yield maturities, and $FTSE_{t-1}$ is a daily, market-based measure of economic uncertainty. Note that we standardize both the revision factors and the yield curve factors. This is helpful for two reasons. First, for a given yield curve factor, we can compare the magnitude of estimated effects for each of our six revision factors. Second, for a given revision factor, we can make relative comparisons of its effect on the three different yield curve factors.[12]

We graph $\psi_{x,j}$ for each yield factor estimated using our full data sample, during expansions, and during contractions in Figure 3.8.[13] Graphs in the left column show estimates of how the U.K.'s yield curve responds to revisions of the BoE's GDP growth forecast. Graphs in the right column do the same for inflation. Within each panel, there are three clusters of three coefficients each surrounded by confidence intervals that fade from 60% (darkest) and to 90% (lightest). Each cluster of coefficients corresponds to either the mean, uncertainty, or skewness revision factor. Within each cluster, we provide estimates from a model that uses revisions of the level, slope, or curvature factor as the dependent variable.

Estimates of Equation (3.3) using our full data indicate that revisions to the higher-order moments of both GDP growth and inflation matter at least as much as revisions of the respective first-order moments. We see that an increase in inflation forecast uncertainty puts upward (downward) pressure on long-term (short-term) maturities and that medium-term yields exhibit a negative response to an increase in outlook on the balance of risks. Also, we see that an increase in the MPC's outlook on the balance of risks for GDP growth can put upward (downward) pressure on

[12]We also show the response of individual yields across the U.k.'s yield curve to information shocks in our appendix.
[13]We also provide coefficient estimates in tabular form in the appendix.

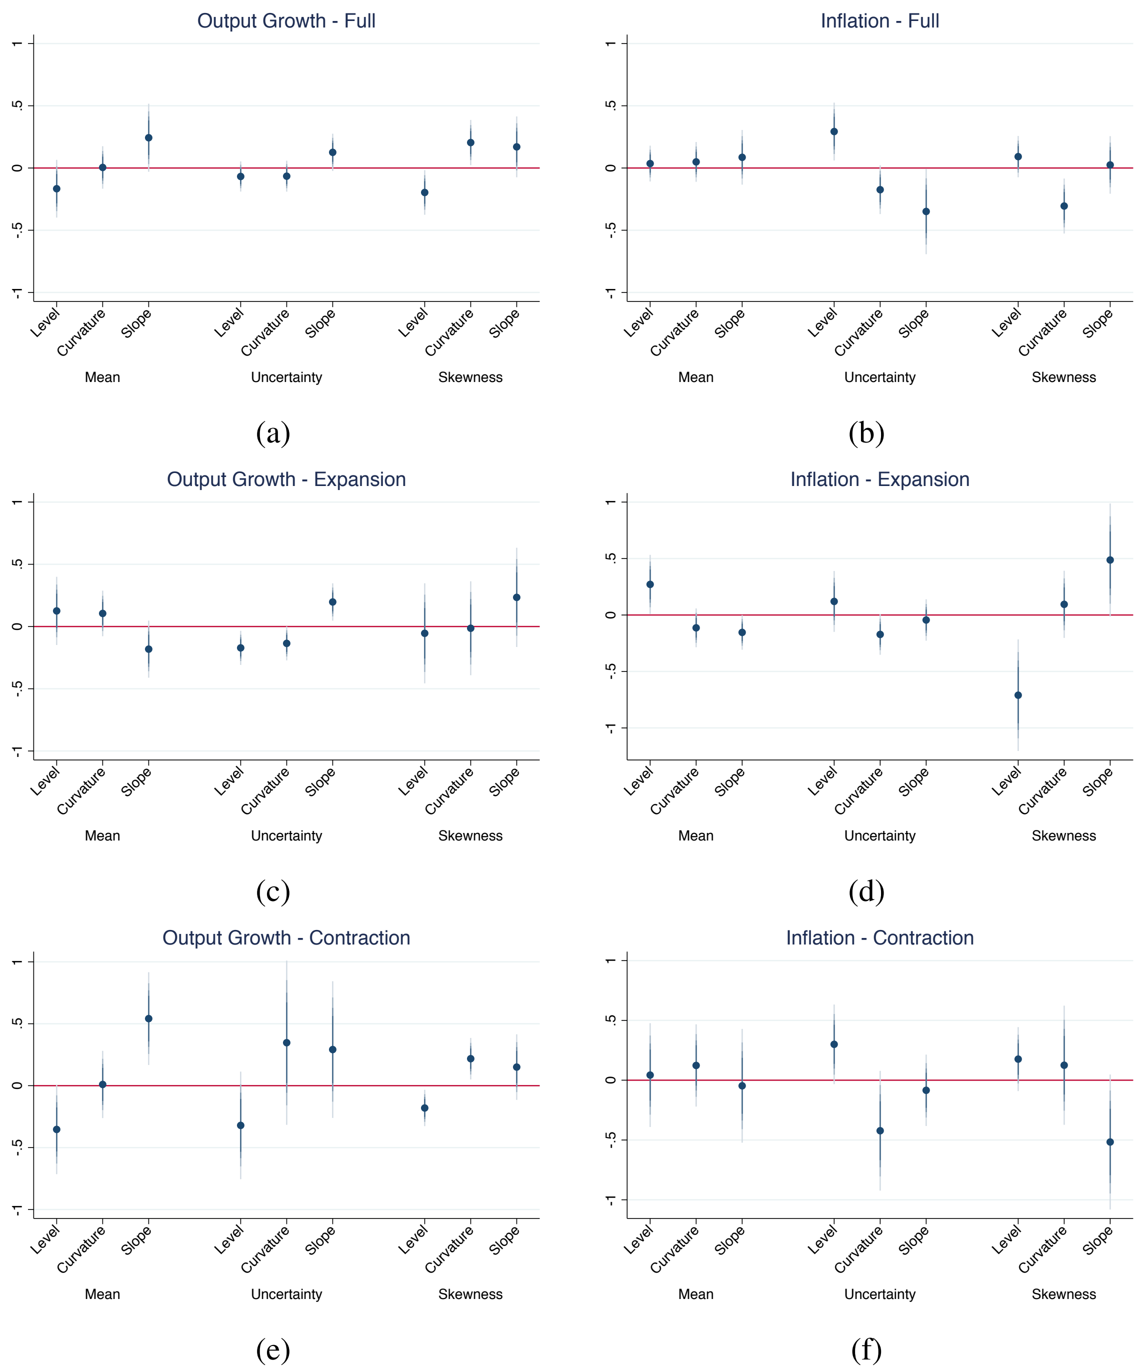

Figure 3.8: Yield responses to information shocks

short- and medium-term (long-term) rates. This suggests that information contained in the third-moment of MPC's predictive density of GDP growth might contribute to yield curve inversion, which is commonly interpreted as a sign of an impending recession.

We gain additional insight by considering the state-contingent effects of information shocks. To do this, we extract the cyclical component of U.K. GDP using the Hamilton filter (Hamilton, 2018) and partition our data into expansionary and contractionary periods whenever the cyclical component of GDP is positive or negative, respectively. We can then consider separately the effects of equivalent information shocks during expansions and contractions. We graph the results of this exercise in panels (c) thru (f) of Figure 3.8.

Panels (c) and (d) of Figure 3.8 depict the results of estimating Equation (3.3) over expansionary periods. For GDP growth, the response of yield factors to first-moment revisions is reversed relative to full-sample estimates. Further, we see that yields exhibit a muted response to skewness shocks during expansions relative to full-sample estimates.

We also find qualitative differences in how yields respond to information shocks surrounding inflation during expansions relative to full-sample estimates. The response of long- and short-term yields to skewness shocks is exacerbated while the response of the yield curve to uncertainty shocks is muted.

First-order moments of the MPC's output growth forecast are much more important during contractions than during expansions, with positive shock putting significant upward (downward) pressure on short-term (long-term) yields. Further, we now see that shocks to both higher-order moments of the GDP growth forecast put upward pressure on short- and medium-term yields, and downward pressure on long-term yields.

Finally, we consider how yields respond to inflation information shocks during a contraction. Most interesting to note is that the yield curve exhibits a qualitatively different response to inflation skewness shocks during contractions relative to expansions.

We ensure the validity of our results by performing a placebo check wherein we re-estimate Equation (3.3) using changes to yield curve factors that occur on days when the BoE did not release a QIR. If our main results reflect a true causal relationship between financial markets and information contained in the higher-order moments of the BoE's forecast, we would expect that our placebo estimates to be relatively precise zeros, with some tolerance for type one errors. We

R^2	Full	Contraction	Expansion
$Level$	0.21	0.43	0.31
$Curvature$	0.26	0.47	0.09
$Slope$	0.22	0.45	0.29

show the results of this placebo test in Figure 3.9.[14] Because we estimate 56 coefficients in this placebo exercise, we could reasonably expect approximately 5 or 6 significant coefficients in this graph. We see fewer than this, which we interpret as some evidence that our main results are not driven by spurious correlation.

We also consider how well our information shocks explain changes in short-, medium-, and long-term yields over our full sample, during contractions, and during expansions. To do this, we compare the coefficient of determination obtained from estimating Equation (3.3) over each possible data sample. We report these results in Section 3.2. Overall, the information shocks better explain yield curve movements during contractions than expansions.

3.2.1 Expectations

This section of the results deals with the relationship between higher-order moments of central bank forecasts and private expectations. Much empirical and experimental work supports the idea that central banks can use point forecasts to coordinate expectations and nudge boundedly-rational agents toward ex-ante rationality, thereby reducing economic volatility. However, there is very limited evidence regarding how higher-order moments influence expectation formation, if at all. We provide some suggestive evidence that professional forecasters do respond to higher-order forecast moments.

To do this, we estimate Equation (3.4) using as our outcome variables the three-, six-, and 12-month-ahead consensus forecast and forecast disagreement[15] for the three-month Sterling rate and the 10-year Guilt rate, avialable at a monthly frequency from Blue Chip Financial Forecast

[14] We provide corresponding numerical coefficient estimates in tabular form in our appendix.

[15] Measured as within-period forecast variation.

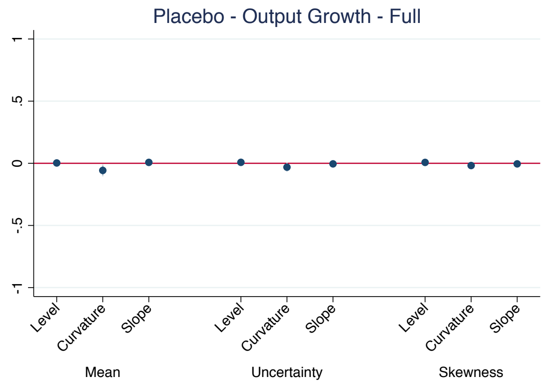

(a)

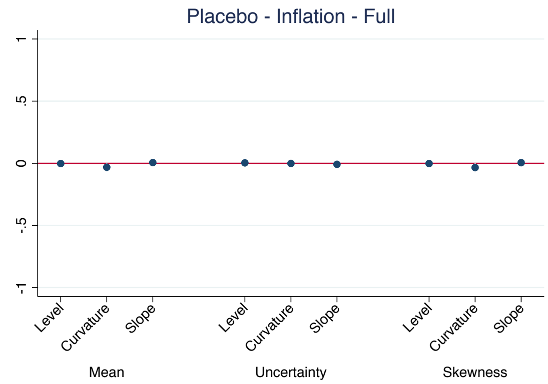

(b)

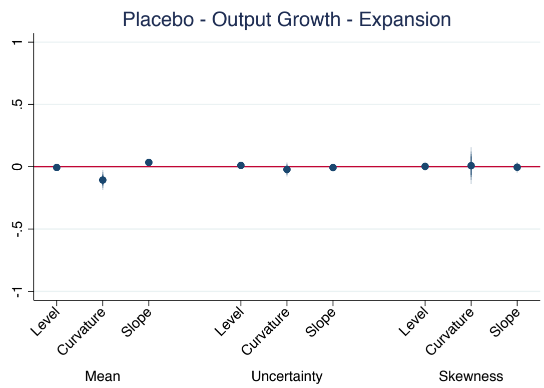

(c)

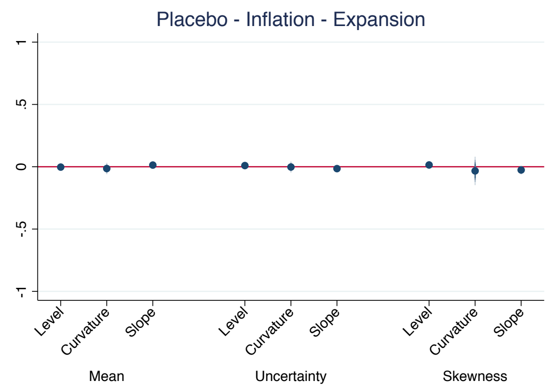

(d)

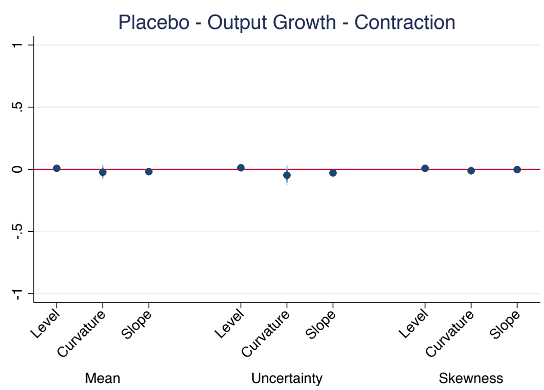

(e)

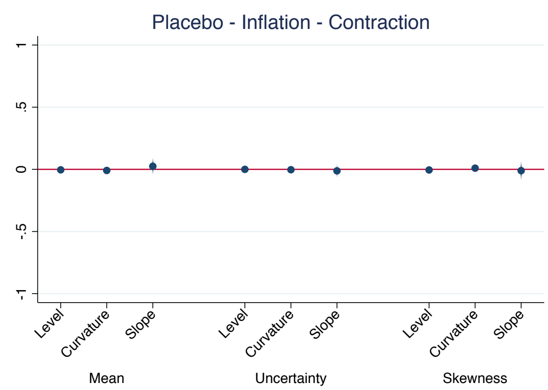

(f)

Figure 3.9: Placebo estimates

(BCFF). The BCFF is a monthly survey administered at the beginning of each month. Because the BCFF is administered early in each month, survey respondents have already provided responses before the QIR release in QIR release months. This means that expectations cannot respond con-

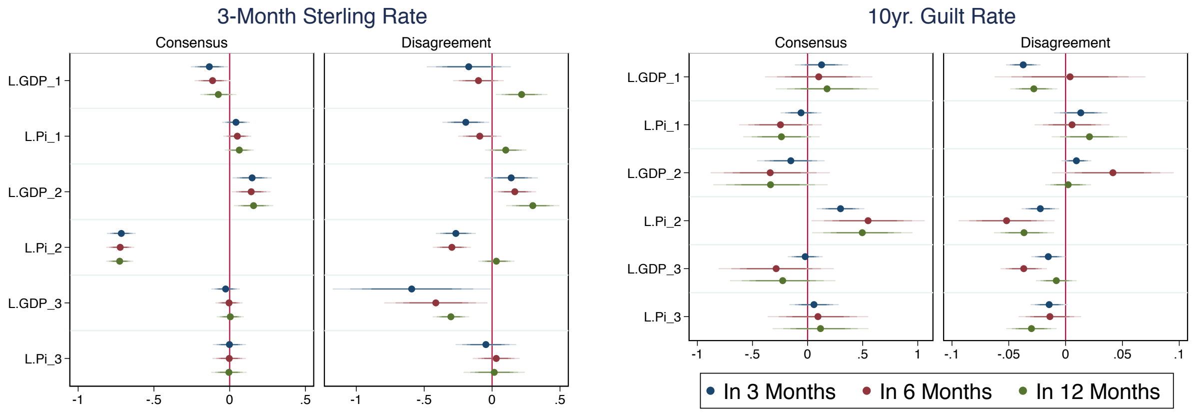

Figure 3.10: Response of professional forecasters to information shocks

temporaneously to information contained in the QIR. We account for this by projecting each of our three forecast outcomes onto lagged values of our information shocks. For example, the BoE releases its first-quarter inflation report about halfway through February each year. To understand how private expectations respond to information in this report, we project our variables of interest from the BCFF survey collected in March onto information shocks extracted from the February report. We graph these results in Figure 3.10.

$$\Delta Measure_{i,t} = \alpha_{i,t} + \sum_{x,j} \psi_{x,j} PC_{x,j,t-1} + \kappa_{i,t} FTSE_{t-1} + \epsilon_{i,t} \tag{3.4}$$

We first consider the relationship between short-term interest rate forecasts and our information shocks. Note that consensus forecasts at all horizons are at least as strongly correlated to forecast uncertainty as they are to first-moment revisions. This correlation is particularly strong for inflation uncertainty, which corresponds to our full-sample estimate of how short-term maturities respond to inflation uncertainty shocks. This logical coherence between expectations and yield movements also holds for GDP growth uncertainty, where both expectations and actual maturities exhibit a small, positive response to output growth uncertainty shocks. However, private expectations of short-term rates are seemingly uncorrelated with the third-moment of either predictive density.

Additionally, we see that inflation uncertainty and forecast disagreement are negatively correlated but GDP growth uncertainty and disagreement are positively correlated. This suggests market participants in the UK may better understand how the MPC, which officially follows an inflation-targeting regime, will respond to inflation uncertainty than to output growth uncertainty. Finally, we also see that expectations of short-term rates are negatively correlated with the third-moment of the MPC's forecast of GDP growth but are seemingly uncorrelated with the third-moment of the inflation forecast. This suggests that market participants better understand how the MPC will respond to upside inflation risk than to upside output risk.

Second, we consider the relationship between long-term interest rate forecasts and information shocks. Similar to short-term rates, we see that private expectations are at least as strongly correlated with higher-order information as they are with information about the central moment of both inflation and GDP growth density forecasts. We also see the same logical coherence between how expectations and maturities respond to information shocks.

Overall, the coherence we observe between how private expectations and maturities respond to second-order information shocks suggests that market participants better understand how maturities respond to uncertainty than to information surrounding risk.

3.2.2 Historical Decomposition

This section considers the historical importance of higher-order information for explaining yield curve changes by decomposing changes in level, slope, and curvature factors over time. We plot these historical decompositions in Figure 3.11.

This decomposition exercise shows that our results are not driven by anomalous events (i.e. the Great Recession). Rather, we see that higher-order information shocks have played a consistently-important role in explaining yield curve movements. Further, we see that higher-order information shocks have historically been at least as important for explaining yield changes as have first-moment information shocks.

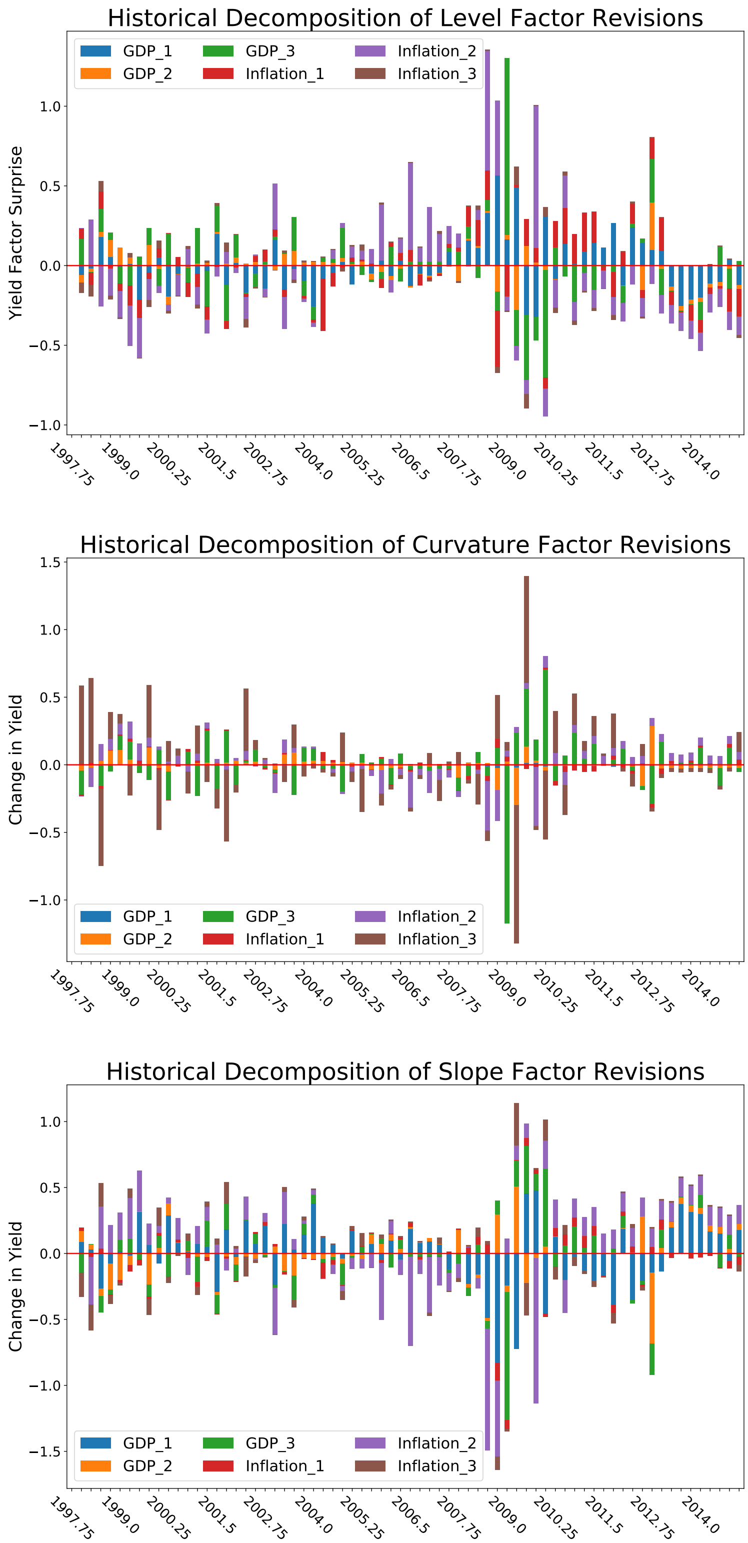

Figure 3.11: Historical decomposition of yield curve factor revisions

3.3 Conclusion

This paper considers the BoE's density forecasts and its revisions to quantify the effects of information flow on the financial markets and expectations. The paper contributes to the broader literature on the effects of news on financial markets and to the emerging literature on how information in the higher-order moments of central bank forecasts matters from a policy perspective.

We find that financial markets respond at least as strongly to the information contained in the higher-order moments of the BoE's density forecasts of output growth and inflation as to the information contained in the corresponding first moments. further, we find that both the magnitude and direction of responses are state-contingent.

Additionally, we use Blue Chip Financial Forecast data to study how professional forecasters respond to information contained in higher-order forecast moments. We find that the consensus forecast and level of forecast disagreement of both short- and long-term interest rates are strongly correlated with higher-order forecast moments. Further, we observe a logical coherence between private expectations and realized yield changes, suggesting that market participants understand how rates will respond to higher-order forecast moments.

Overall, our results suggest that communicating high-order forecasts moments to market participants does affect their subsequent behavior. Important from a policy perspective is that higher-order moments can move markets, which suggests that density forecasting is a viable policy option.

3.4 References

1. Ahrens, Steffen, Joep Lustenhouwer, and Michele Tettamanzi. "The Stabilizing Effects of Publishing Strategic Central Bank Projections."

2. Andrade, Philippe, and Filippo Ferroni. “Delphic and Odyssean monetary policy shocks: Evidence from the euro area." Journal of Monetary Economics (2020).

3. Arifovic, Jasmina, and Luba Petersen. “Stabilizing expectations at the zero lower bound: Experimental evidence." Journal of Economic Dynamics and Control 82 (2017): 21-43.

4. Blinder, Alan S., et al. “Central bank communication and monetary policy: A survey of theory and evidence." Journal of Economic Literature 46.4 (2008): 910-45.

5. Brand, Claus, Daniel Buncic, and Jarkko Turunen. “The impact of ECB monetary policy decisions and communication on the yield curve." Journal of the European Economic Association 8.6 (2010): 1266-1298.

6. Dahlhaus, Tatjana, and Tatevik Sekhposyan. “Monetary policy uncertainty: A tale of two tails." No. 2018-50. Bank of Canada, 2018.

7. Diebold, Francis X., and Canlin Li. “Forecasting the term structure of government bond yields." Journal of econometrics 130.2 (2006): 337-364.

8. Elder, Robert. "Assessing the MPC’s fan charts." Bank of England Quarterly Bulletin, Autumn (2005).

9. Eusepi, Stefano, and Bruce Preston. “Central Bank Communication and Expectations Stabilization." American Economic Journal: Macroeconomics, 2 (3): 235-71 (2010).

10. Gürkaynak, Refet S. “Using federal funds futures contracts for monetary policy analysis." (2005).

11. Gürkaynak, Refet S., Brian Sack and Eric T. Swanson. “The Excess Sensitivity of Long Term Interest Rates: Evidence and Implications for Macroeconomic Models." American Economic Review, 95, pp.425-436 (2005a).

12. Gürkaynak, Refet S., Brian P. Sack, and Eric T. Swanson. “Do actions speak louder than words? The response of asset prices to monetary policy actions and statements." (2004).

13. Hamilton, J. D. (2018). Why you should never use the Hodrick-Prescott filter. Review of Economics and Statistics, 100(5), 831-843.

14. Hansen, Stephen, Michael McMahon, and Matthew Tong. "The long-run information effect of central bank communication." Journal of Monetary Economics 108 (2019): 185-202.

15. Hoesch, Lukas, Barbara Rossi, and Tatevik Sekhposyan. "Has the information channel of monetary policy disappeared? Revisiting the empirical evidence." (2020).

16. Hubert, Paul. "FOMC forecasts as a focal point for private expectations." Journal of Money, Credit and Banking 46.7 (2014): 1381-1420.

17. Hubert, Paul. "The influence and policy signalling role of fomc forecasts." Oxford Bulletin of Economics and Statistics 77.5 (2015): 655-680.

18. Hubert, Paul. State-dependent effects of monetary policy: the central bank information channel. OFCE, 2019.

19. Maule, Becky, and Paul Hubert. "Policy and macro signals as inputs to inflation expectation formation." (2020).

20. Kryvtsov, Oleksiy, and Luba Petersen. “Central bank communication that works: Lessons from lab experiments." Journal of Monetary Economics (2020).

21. Miranda-Agrippino, Silvia, and Giovanni Ricco. "The transmission of monetary policy shocks." (2018).

22. Mokhtarzadeh, Fatemeh, and Luba Petersen. “Coordinating expectations through central bank projections." Available at SSRN 2915814 (2019).

23. Morris, Stephen, and Hyun Song Shin. “Social value of public information." american economic review 92.5 (2002): 1521-1534.

24. Nakamura, Emi, and Jón Steinsson. "High-frequency identification of monetary non-neutrality: the information effect." The Quarterly Journal of Economics 133.3 (2018): 1283-1330.

25. Petersen, Luba and Rholes, Ryan. "Does Policy Communication Matter More During a Pandemic?" (2020).

26. Rholes, Ryan, and Luba Petersen. “Should central banks communicate uncertainty in their projections?" Available at SSRN 3529832 (2020).

27. Rudebusch, Glenn D., and John C. Williams. “Revealing the secrets of the temple: The value of publishing central bank interest rate projections." Asset Prices and Monetary Policy. University of Chicago Press, 2008. 247-289.

28. Swanson, Eric T. “Have increases in Federal Reserve transparency improved private sector interest rate forecasts?." Journal of Money, Credit and Banking (2006): 791-819.

29. Woodford, Michael. “Central bank communication and policy effectiveness." No. w11898. National Bureau of Economic Research, 2005.

30. Wu, Jing Cynthia, and Fan Dora Xia. “Measuring the macroeconomic impact of monetary policy at the zero lower bound." Journal of Money, Credit and Banking 48.2-3 (2016): 253-291.

4. SHOULD CENTRAL BANKS COMMUNICATE UNCERTAINTY IN THEIR PROJECTIONS?*

[1]

4.1 Introduction

Central banks have become increasingly transparent over the last few decades, with most banks now disclosing information surrounding operations, procedures, economic outlook and policy. This transparency revolution is driven largely by a deeper understanding of the importance of expectations, central bank credibility, and of the ability of communication to function as a policy tool.

A prominent feature of transparency is the publication of macroeconomic projections. As argued by Greenspan (2004), forming and communicating macroeconomic projections plays as an important role in the preemptive response of policy makers to inflationary pressures. Such forecasts not only play an important internal role in policy deliberations, but also provide market participants with insight into how the central bank thinks the economy and policy rate may evolve. Projections align private-sector expectations and improve forecast accuracy in theory (Geraats (2002), Woodford (2005), Rudebusch and Williams (2008), Gosselin et al. (2008), Blinder et al. (2008)), experiments (Kryvtsov and Petersen, 2015, 2020, Mokhtarzadeh and Petersen 2018, Ahrens et al. 2019), and in practice (Brubakk et. al. (2017), Hubert (2014, 2015), Blinder et al. (2008), and Kool and Thornton (2015)).

However, central banks are communicating in an uncertain world. Not only are the timing and magnitude of the effects of monetary policy uncertain, but so are the shocks the economy faces. Consequently, many central banks publish density forecasts, rather than just point projections, in an effort to convey a subjective measure of uncertainty about economic outlook and the future path

[1] *Reprinted with permission from "Should central banks communicate uncertainty in their projections?" by Rholes, Ryan and Petersen, Luba, 2021. Journal of Economic Behavior and Organization, Volume 183, Pages 320-341, Copyright [2021] by Elsevier.

of policy and to preserve credibility. Density forecasts typically convey the same information contained in point forecasts, but also present the central bank's uncertainty surrounding its projections (second moment) and the bank's outlook on risk (third moment). The Bank of England was the first to publish 'fan charts' of its macroeconomic projections in 1998, with the Federal Reserve, the European Central Bank, the Reserve Bank of Australia, the Bank of Canada, and the Swedish Riksbank following suit.

Despite the growing trend of central banks communicating uncertainty by publishing density forecasts, there exists almost no empirical or theoretical evidence that this improves the ability of central bank projections to influence markets, or coordinate and improve private forecasts. One exception is Rholes and Sekhposyan (2020) who show that short-term yields respond at least as strongly to revisions of the second- and third-moments of the BoE's density forecasts as they do to revisions of the first-moment of the same density forecasts.[2] In closely related experimental work, Mokhtarzadeh and Petersen (2018) show that density projections (that present both the point forecast and a confidence interval) are effective at managing expectations if they are relevant and easy to understand. Their findings, however, do not isolate the effect of density forecasts from the point projections.

This paper provides original empirical evidence on the effects of point and density forecasts on the management and formation of inflation expectations. We systematically vary inflation projection announcements communicated by the economy's automated central bank within a macroeconomic learning-to-forecast laboratory experiment where groups of participants simultaneously form inflation expectations. We incentivize participants to form accurate one- and two-period-ahead inflation expectations. Aggregated expectations endogenously influence macroeconomic dynamics. Given participants' potentially bounded rationality, there is a role for central bank communication to guide expectations. We also elicit participants' confidence about their forecasts, allowing us to clearly identify the transmission of central bank uncertainty to forecasters.

We consider three levels of central bank communication in a between-subject design: No sup-

[2]Uncertainty about monetary policy can have negative economic effects. See Neely (2005), Swanson (2006), Bauer (2012), Husted, Rogers, and Sun (2018, 2019) for discussion.

plementary communication, five-period ahead point projections, and five-period ahead point and density projections. Both projections are based on the assumption that agents form ex-ante rational expectations. Density projections are symmetric one-standard deviation confidence intervals around the point projection. This variation allows us to disentangle the effects of communicated uncertainty on expectation formation.

Relative to a baseline of no communication, we find that point projections reduce disagreement and uncertainty about future inflation, and medium term (two-period-ahead) forecast errors. Moreover, point projections increase the proportion of inexperienced participants who forecast one-period-ahead inflation as if they were ex-ante rational by 72 percentage points for a total of 86% of participants.

Density projections mute the beneficial effects of point projections. Compared to point projections, communicating density forecasts significantly increase forecast errors, uncertainty, and disagreement about two-period-ahead inflation. Credibility in the central bank's point projection is significantly lower when it includes a density forecast. Only 57% of inexperienced participants in density treatments form rational one-period-ahead expectations.

It is also important to understand why projections have proven effective at managing real-word expectations. Is it because economic agents use publicly communicated projections purely as a coordination device or do the projections provide valuable information that forecasters and market participants use to improve forecast accuracy? To answer this question, we conduct the same communication treatments in an individual-choice environment absent any strategic considerations. In both Individual and Group treatments, the projection provides information and, more importantly, reduces the complexity of the forecasting problem. In the Group treatments, there is an additional strategic consideration. Group participants should use the projection if and only if they believe the majority of participants will.

Thus, our individual-choice treatments have participants play the role of the representative forecaster, with their own expectations employed as the aggregate expectation driving macroeconomic dynamics. Our motivation for this treatment is to investigate whether people choose to use projections because of their information content or as a coordination device. Thus, our individual-forecaster treatment eliminates any coordination motives. Though this individual setting lacks some degree of external validity, it allows us to tease apart and understand the motives underlying central bank projections. We expose participants in our individual-choice treatments to the same three levels of central bank communication used in our group setting. This allows us to draw inference about the effect of strategic motives on how subjects use central bank forecasts when forming expectations. To the best of our knowledge, this is the first learning-to-forecast experiment to compare individual vs. group forecasting behaviour.

Absent strategic motives, participants are significantly more heterogeneous in their forecasts and form larger forecast errors. Individual forecasters also anticipate making larger forecast errors when they have no supplementary communication from the central bank, suggesting that the wisdom of the group improves confidence. Point projections reduce individuals' two-period-ahead forecast errors, though not as effectively as in group settings. Neither point projections nor density projections consistently reduce disagreement or uncertainty in Individual treatments. Our findings suggests that the information content associated with projections is not as valuable as their ability to serve as a coordination device.

Finally, our experimental results provide useful guidance for the modeling of inflation expectations. First, we find ample evidence to suggest that a large majority of participants will adopt an as-if rational heuristic when they observe a rationally-constructed inflation point projection. Second, most participants use the same heuristics to formulate both their short and medium term expectations.

4.2 Experimental Design

Our experiment seeks to understand how central bank point and density projections influence expectations and aggregate dynamics. To this end, we study a 'learning-to-forecast (LTF)' experimental macroeconomy that uses either the aggregate expectations of groups or the expectations of individuals, depending on the treatment, to influence aggregate variables. Such experimental economies are well-studied. Macroeconomists have used similar experiments to study expectation formation and equilibria selection (Adam, 2007), the effects of different monetary policy rules and targets on expectation formation (Pfajfar and Žakelj 2014, 2018; Assenza et al. 2013, Hommes et al. 2019a; Cornand and M'Baye, 2018a), expectation formation at the zero lower bound (Arifovic and Petersen 2017, Hommes et al. 2019b), and the endogenous dynamics of expectations and real decisions (Bao et al., 2013). Learning-to-forecast experiments have been shown to reasonably match inflation forecasting patterns observed in surveys of households, firms, and professional forecasters (Cornand and Hubert, 2019).

We are also interested in understanding how subjects' own uncertainty about future inflation responds to both precise and noisy central bank projections. Pfajfar and Žakelj (2016) also explore uncertainty in response to different inflation targeting regimes. Similarly, our paper is closely related to Cornand and M'baye (2018b), in which the authors use an LTF framework to explore the relative merits of point and band inflation targeting. The authors find that during periods of low uncertainty, band targeting better stabilizes inflation. Also, point targeting with tolerance bands leads to a lower and less volatile output gap and interest rate. The authors also find that both regimes are equally ineffective at stabilizing macroeconomic dynamics during periods of high uncertainty. Our paper differs from theirs in that, rather than varying the central bank's targeting regime, we use a single point inflation target and vary the central bank's communication strategy.

We begin by describing the design of our baseline environment, which involves groups of seven participants playing the roles of forecasters in an environment with no supplementary central bank communication. In Section 4.2.3, we describe how the environment changes as we allow for individually-driven economies and central bank projections.

We summarize the flow of information, decisions, and outcomes throughout the experiment in Figure 4.1. Each experiment consists of two different sequences of 30 sequentially linked periods. In each period $t \in [1, 30]$, participants submit forecasts about $t + 1$ and $t + 2$ inflation, as well as predictions about the magnitude of their forecast errors.

At the beginning of each period, subjects observe all historical information about inflation, the nominal interest rate, and demand shocks. Importantly, subjects can also observe the value of the current-period demand shock.[3] Subjects also observe their own history of one- and two-period-ahead inflation forecasts and total earnings.

Figure 4.1: Flow of information, decisions, and outcomes

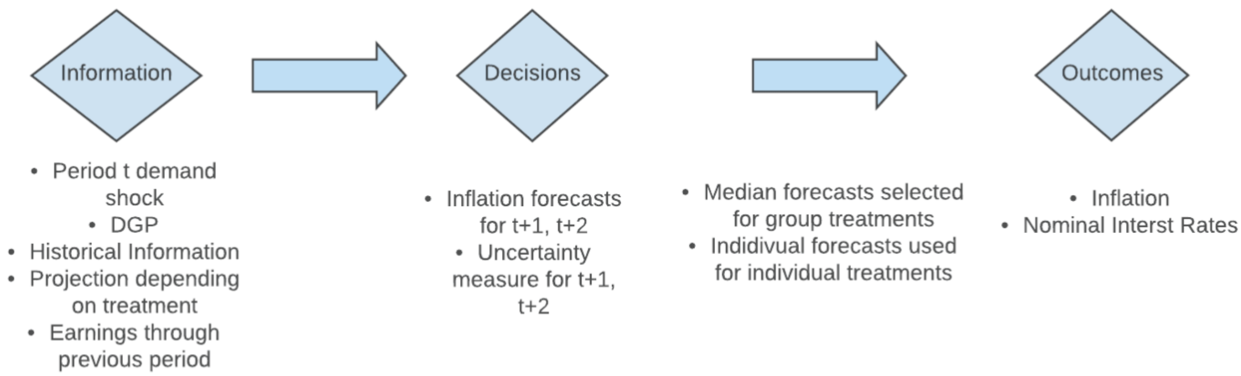

Subjects had 65 seconds to form forecasts in periods 1-9 of each sequence and 50 seconds thereafter. Inflation expectations and corresponding uncertainty forecasts were submitted in basis points. Inflation forecasts could be positive, negative, or zero. Uncertainty measures could be either zero or positive. All submissions were unbounded. Since we collect forecast in terms of basis points, subjects could submit forecasts with a precision of $\frac{1}{100}$th of 1%.

After all subjects submitted expectations or time elapsed, participants moved onto the next period. The economy's data-generating process, which will be described in the next subsection, relies on aggregate one- and two-period-ahead inflation expectations. We use the median forecasts, rather than the average, as the aggregate expectations to curtail the impact that any one subject can have on our experimental economies. This has the effect of making it as though forecasters are atomistic.

[3]Subjects have sufficient information to calculate the expected value of future shocks and can incorporate this, if they desire, into current-period forecasts.

4.2.1 Data-generating process

Each treatment shares a common data-generating process, which is derived from a log-linearized, representative-agent New Keynesian (NK) framework We re-write this model to eliminate the need for expectations about the one-period-ahead output gap. This manipulation of the NK model allows us to use a system of equations driven by one- and two-period-ahead inflation expectations and aggregate disturbances. Thus, we need only elicit $\mathbb{E}_t\{\pi_{t+1}\}, \mathbb{E}_t\{\pi_{t+2}\}$ from our subjects. See the Online Appendix Section B for details about the differences in stability and forecast errors in the two formalizations.

We begin with a standard 4-equation, reduced-form NK model

$$\pi_t = \beta\mathbb{E}_t\{\pi_{t+1}\} + \kappa x_t \tag{4.1}$$

$$i_t = \phi_\pi \pi_t + \phi_x x_t \tag{4.2}$$

$$x_t = \mathbb{E}_t\{x_{t+1}\} - \sigma^{-1}[i_t - \mathbb{E}_t\{\pi_{t+1}\} - r_t^n]. \tag{4.3}$$

We eliminate the need to elicit $\mathbb{E}_t\{x_{t+1}\}$ by rewriting (3), iterating forward, taking expectations, applying the law of iterated expectations, and substituting to obtain:

$$x_t = (\kappa^{-1} + \sigma^{-1})\mathbb{E}_t\{\pi_{t+1}\} - \beta\kappa^{-1}\mathbb{E}_t\{\pi_{t+2}\} - \sigma^{-1}i_t + \sigma^{-1}r_t^n. \tag{4.4}$$

Substitutions yield a representation of (3) that depends only on inflation expectations,

$$\pi_t = [\beta + \kappa\gamma_1\gamma_2]\,\mathbb{E}_t\{\pi_{t+1}\} - \gamma_1\beta\mathbb{E}_t\{\pi_{t+2}\} + \kappa\gamma_1\sigma^{-1}r_t^n, \tag{4.5}$$

where we use the following variable substitutions

$$\gamma_1 = \left(\frac{\sigma + \phi_\pi\kappa + \phi_x}{\sigma}\right)^{-1} \tag{4.6}$$

$$\gamma_2 = (\kappa^{-1} + \sigma^{-1} - \sigma^{-1}\phi_\pi\beta). \tag{4.7}$$

This yields a dynamical system that can be solved using $\mathbb{E}_t\{\pi_{t+1}\}, \mathbb{E}_t\{\pi_{t+2}\}$, r_t^n. Here, r_t^n represents a demand shock that evolves following an AR(1) process,

$$r_t^n = \rho_r r_{t-1} + \epsilon_{r,t} \tag{4.8}$$

where $\epsilon_{r,t}$ is i.i.d. $\sim \mathcal{N}(0, \sigma_r)$ and ρ_r is a persistence parameter. The data-generating process is calibrated to match moments of Canadian data following Kryvtsov and Petersen (2015); $\sigma = 1$, $\beta = 0.989$, $\kappa = 0.13$, $\phi_\pi = 1.5$, $\phi_x = 0.5$, $\rho_r = 0.57$, and $\sigma_r = 113$ bps, with a steady state of $\pi^* = x^* = 0$

Given these parameters, the system of equations reduces to

$$\pi_t = 1.54\mathbb{E}_t\{\pi_{t+1}\} - 0.58\mathbb{E}_t\{\pi_{t+2}\} + 0.08r_t^n \tag{4.9}$$

$$i_t = 4.44\mathbb{E}_t\{\pi_{t+1}\} - 3.12\mathbb{E}_t\{\pi_{t+2}\} + 0.41r_t^n. \tag{4.10}$$

We use aggregate expectations provided by participants to close the model. Aside from Adam (2007), this is the only experiment within a NK framework that elicits expectations for two future time periods. However, this particular formulation of the NK model is novel to the learning-to-forecast literature. This formulation accomplishes two things. First, it reduces the cognitive complexity subjects face by allowing them to focus on forecasting a single time series. Second, it allows us to understand how these information conditions affect expectations further into the future than would be possible otherwise.

Worth noting here is the counter-balancing effect of expectations on this system. Equation (4.9) and Equation (4.10) retain the familiar feature that one-period-ahead expectations are self-fulfilling, but we also observe, counter-intuitively, that two-period-ahead expectations are not self-fulfilling. However, this counter-balancing of expectations makes sense from the perspective of consumption smoothing. Expecting higher prices tomorrow encourages an agent to consume more today to avoid the higher prices tomorrow. This behavior puts upward pressure on prices today, leading to higher inflation today. If an agent also expects inflation two days from now, then they will want to have more money to spend tomorrow than otherwise, so that an agent can similarly avoid paying higher prices two days from now. As we show in the Online Appendix, this particular presentation of the DGP does not alter the qualitative benefits of rationally–constructed central bank projections.

4.2.2 Payoffs

We incentivized forecasts using the scoring rule described by Equation (4.11). Notice that F_{it} exhibits exponential decay as that forecaster $i's$ absolute forecasting error increases.

$$F_{it} = 2^{-.5|\mathbb{E}_{i,t-1}\{\pi_t\}-\pi_t|} + 2^{-.5|\mathbb{E}_{i,t-2}\{\pi_t\}-\pi_t|} \tag{4.11}$$

Subjects received payoffs for all forecasts about $t+1$ formed in $t \in [1, 29]$ and $t+2$ forecasts formed in $t \in [1, 28]$. Subjects in our experiment also provided measures of uncertainty about their one- and two-period-ahead inflation forecasts, which we denote here as $u_{i,t+1}, u_{i,t+2} \in \mathbb{N}_0$. This measure of uncertainty creates a subject-level density forecast in each period for both forecast horizons. We assume a subject's forecast uncertainty is symmetric around her point forecasts, which is similar to our assumption about the central bank's forecast uncertainty. We incentive this uncertainty measure using a piece-wise scoring rule.[4] A subject earns nothing if actual inflation fall outside her density forecast (i.e., her uncertainty bands). Otherwise, a subject earns $U_{i,t+k}$, where $k = \{1, 2\}$:

$$U_{i,t+k} = \frac{15}{10 + u_{i,t+k}} \tag{4.12}$$

The payoff that participants receive for their error forecast decreases in the level of their forecasted error. Because we incentivize uncertainty measures for each forecast horizon, a subject

[4] A concern here is that this scoring rule may only be incentive-compatible with risk-neutral agents. A risk-loving agent may slightly under-report her uncertainty while a risk-averse agent may slightly over-report uncertainty. However, we can distinguish neither risk-loving behavior from overconfidence nor risk-averse behavior from under-confidence.

could earn a total of three points in each period for her uncertainty measures. This scoring rule is similar to the rule used in Pfajfar and Žakelj (2016), which studied the effect of various monetary policy rules on individual uncertainty.[5] To address the possibility of hedging, we randomly selected at the session level in each period whether to pay F or $U_{t+1} + U_{t+2}$. However, we never paid both in the same period.

4.2.3 Treatments

We used a 3x2 between-subject experimental design to study the effects of central bank communication and strategic motives on expectation formation and aggregate dynamics. Table 4.1 summarizes the treatments.

Subjects formed forecasts under one of three information conditions: a baseline where a mechanistic central bank provided no projections (NoComm), a projection-only treatment where the central bank provided an evolving five-period-ahead point forecast of inflation (Point), and a density forecast treatment where the central bank provided both an evolving five-period-ahead point forecast point and density forecast of inflation (Point&Density).

We also varied the environment along a coordination dimension. Subjects either participated in a Group treatment, where they interacted in an experimental economy with six other subjects, or an Individual treatment, where each subject served as the sole forecaster in her experimental economy. In other words, subjects in Individual treatments played an individual choice game; their expectations alone, coupled with the demand shock, drove the dynamics of their economies. Subjects in Individual treatments understood that they each inhabited a unique economy.

Participants interacted in an online platform. See **??** for an example of the NoComm interface, **??** for the Point interface, and **??** for the Point&Density interface. Subjects in all treatments always interacted with the same screen in each decision period. The screen updated to display new information as that information became available.

Aside from the communications from the central bank, all participants received common information. The top left corner of a subject's screen showed the subject's identification number, the

[5] Pfajfar and Žakelj (2016) elicit participants' own 95% confidence intervals around their point forecasts.

Table 4.1: Treatments summary

Group				
CB Projection	**Sequences**	**Total Subjects**	**Periods**	**Aggregate Expectations**
NoComm	6	42	30 x 2	median of group
Point	6	42	30 x 2	median of group
PointDensity	6	42	30 x 2	median of group
Individual				
CB Projection	**Sequences**	**Total Subjects**	**Periods**	**Aggregate Expectations**
NoComm	6	39	30 x 2	own
Point	6	42	30 x 2	own
Point&Density	6	38	30 x 2	own

current decision period, time remaining to make a decision, and the total number of points earned through the end of the previous period. The interface also featured three history plots. The top history panel displayed past interest rates, and both past and current shocks. The second panel displayed the subject's one-period-ahead inflation forecast (blue dots), the subject's uncertainty surrounding this one-period-ahead forecast (blue shading), and all realized values of inflation (red dots). The third history panel displayed the subject's two-period-ahead inflation forecast (orange dots), the subject's uncertainty surrounding this two-period-ahead forecast (orange shading), and all realized values of inflation (red dots).

Treatment variation appeared in the second and third history panels. Notice in **??** (NoComm) that the central bank provided neither point nor density forecasts. In **??** (Point) the second and third history plots displayed the central bank's evolving, five-period-ahead point forecast (green dots). Finally, the second and third history plots in **??** (Point&Density) contained the central bank's evolving five-period-ahead point forecast (green dots) with its corresponding level of uncertainty (green shading).

We explained to subjects in both Point and Point&Density treatments that the central bank's projections were not a guarantee, thus indicating that there is some level of uncertainty surrounding the central bank's projections. We also explained to subjects that the central bank's forecasts are based on the DGP and all available information to emphasize the potential errors in the central bank's projection model.

We further indicate in Point&Density treatments that the density forecast represents the central bank's own uncertainty about its point projections. Our exact phrasing was: "These forecasts also include green shading, which represents the Central Bank's level of uncertainty for its corresponding point projections. These bands will contain the correct realization of inflation about 66% of the time."

The level of fundamental uncertainty is always the same because it is driven by the shock process, which is quantitatively conveyed to participants in the instructions. The density treatment does not provide any new information in this regard. It only changes the salience of the uncertainty in the environment.

The mechanistic central bank in our experiment always used a symmetric density forecast. However, this is not always true of density forecasts provided by real-world central banks. An interesting extension to this project would be studying how forecasters react to asymmetric density forecasts. This is akin to studying how forecasters incorporate information contained in the skewness (third central moment) of a central bank density forecasts, which we can think of intuitively as a bank's outlook on economic risks. Finally, we note that the mechanistic central bank assumed that the aggregate expectation in each experimental economy was ex-ante rational. The central bank's density forecast was simply a one standard-deviation band centered around its point forecast.[6]

We conducted six independent sessions of each of the six treatments for a total of 36 experimental sessions. Each session consisted of two, 30-period repetitions (decision blocks). For a given session, we randomly drew a shock sequence for two repetitions from our theoretical distribution of shocks. Within a treatment, we drew six different shock sequences (or 12 if you count each repetition). We then used the same shock sequences across treatments for comparability. Usage of different shock sequences allows for a more robust analysis of expectation formation.

[6]Mokhtarzadeh and Petersen (2020) show experimentally that both the assumptions underlying central bank projections and the information communicated alongside projects matter for expectations formations and aggregate stability. Thus, it seems reasonable that additional information about perceived risk could change how agents incorporate information contained in other forecast moments.

Finally, we initialized each experimental economy at the zero–inflation steady state. We showed subjects five preceding periods of the economy operating along this steady state before introducing a demand shock in period one. Note that it would not be rational, given that shocks evolve according to an autoregressive process, to forecast zero inflation for periods two or three.

4.2.4 Procedures

The experiments were conducted at Simon Fraser University from October 2019 to January 2020. We began each session by reading aloud from paper instructions that included detailed information about subjects' forecasting task, the uncertainty measurement task, how we incentivize forecasts and uncertainty, and how the experimental economy evolved in response to expectations and aggregate shocks. Participants knew they could use the computer's calculator or spreadsheets if desired.

Following the instructions, subjects played four unpaid practice periods during which they could ask questions. Following the practice periods, subjects played through the two incentivized sequences. Each sequence employed a different variation of the shock sequence so that subjects did not repeat an identical game in the second block of decisions. We paid subjects in cash immediately following each experimental session. Payoffs, including a CDN$7 show–up fee, ranged from CDN$12-25, and were on average CDN$21.

4.2.5 Hypotheses

Our experimental design allows us to test several hypotheses regarding differences in how point and density forecasts impact aggregate dynamics and individual behavior. Further, we are able to test hypotheses regarding how subjects use information differently when they face the strategic considerations present in a coordination setting.

If all subjects in our experiment are model–consistent and have full information about the data–generating process (as they do), then we should observe that neither projection (Point or Point&Density) changes aggregate dynamics or individual behavior relative to one another or the NoComm setting. The projections would be irrelevant because they would neither increase the information set of agents (as the mechanistic central bank does not have more information than our subjects) nor impact how agents use available information. Further, model-consistent subjects would behave equivalently when forming expectations in both individual-choice and coordination settings since they possess the same information.

However, ample laboratory and survey evidence demonstrate that individuals, rather than conforming to rational expectations, form expectations in a backward-looking manner (See Assenza et al., 2013; Pfajfar and Santoro, 2010; Pfajfar and Žakelj 2014; Coibion and ˘Gorodnichenko, 2015; Malmandier and Nagel, 2016). In a closely related experiment, Mokhtarzadeh and Petersen (2020) show that, in the absence of central bank forecasts, most subjects follow either a Constant Gain Learning model or an Adaptive(1) heuristic, whereby they equally weight historical information and the ex-ante rational forecast when forming expectations. Backward-looking heuristics also explain behavior in Pfajfar and Žakelj (2016), Cornand and M'Baye (2018b), and Hommes et al. (2019).

Further, extensive empirical and experimental evidence supports the notion that central bank forecasts have effectively coordinated private-sector expectations and stabilized markets (Hubert 2014, 2015, Jain and Sutherland, 2018, Mokhtarzadeh and Petersen, 2018, Ahrens et al. 2018, Kryvtsov and Petersen, 2020).

Figure 4.2 presents simulated standard deviations of inflation and mean absolute forecast errors under various forecasting heuristics and parameterizations. See Table 4.7 for the underlying models and the Online Appendix Section B for a comparison of statistics under the one-period and two-period ahead version of the linearized New Keynesian model under alternative heuristics.

Simulations show that inflation variability and forecast errors are strongly positively correlated. Variability and errors are high when the representative agent adopts simple heuristics such as trend-chasing, Adaptive(1), naive, and constant gain learning. Under all these heuristics, the agent uses a misspecified model that relies on historical data, which introduces persistence into the economy. Further, these heuristics yield relatively unstable dynamics because ad-hoc monetary policy is

Figure 4.2: Simulated statistics under alternative forecasting heuristics

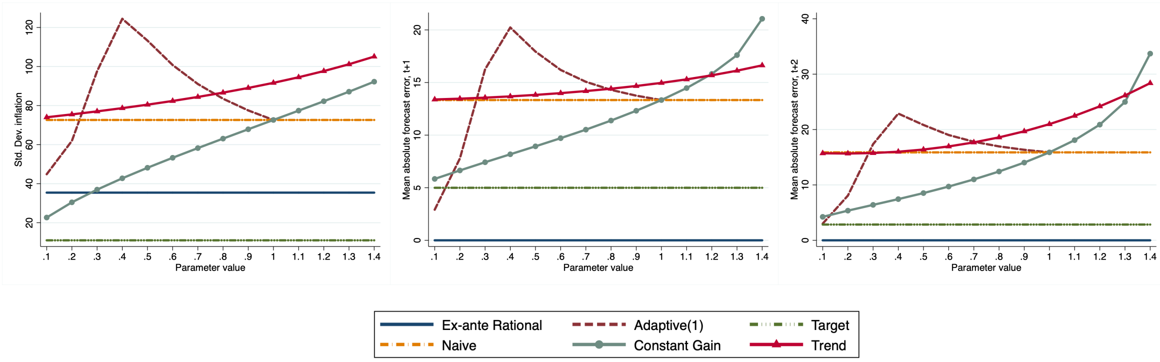

unable to compensate for the agent's non-rational expectations.[7] We expect simple heuristics such as naive, trend-chasing, and constant gain learning models to dominate given past experimental findings in similar experiments (Cornand and M'Baye, 2018a; Pfajfar and Zakelj, 2014 and 2016; Petersen and Rholes, 2020).

While forecasting the inflation target for one- and two-period-ahead inflation would produce the lowest volatility and the smallest forecast errors of any of the non-rational alternatives we consider, we believe it is unlikely that subjects will adopt this heuristic. This is because participants have full information about the economy's data-generating process and know that monetary policy is unable to perfectly offset the exogenous demand shocks.

Heuristics that employ some degree of rational expectations (e.g. Adaptive(1) and ex-ante rational) are considerably more sophisticated and unlikely to be used by most subjects. However, both inflation variability and forecast errors are substantially lower if agents form ex-ante rational expectations. Thus, there is value in easy-to-understand central bank communication that can guide boundedly-rational agents to forecast rationally.

To illustrate how projections can manage expectations and uncertainty, we consider the beliefs of a Bayesian-updating agent. We suppose there is some true state of the world π_{t+s} and that both our economic agent i and our central bank form beliefs about π_{t+s} equal to the state of the world

[7] Note that under optimal monetary policy, the central bank would be able to offset the demand shock and perfectly stabilize inflation given the agent's expectations.

plus some error,

$$\mathcal{B}_i \equiv \mathbb{B}_{i,t}\{\pi_{t+s}\} = \pi_{t+s} + \delta_{i,t}, \;\; \delta_{i,t} \sim \mathcal{N}\left(0, \frac{1}{\psi}\right), \tag{4.13}$$

$$\mathcal{B}_{CB} \equiv \mathbb{B}_{CB,t}\{\pi_{t+s}\} = \pi_{t+s} + \nu_t, \;\; \nu_t \sim \mathcal{N}\left(0, \frac{1}{\omega}\right). \tag{4.14}$$

$\mathcal{B}_i$ is the ex-ante expectation of agent i about π_{t+s} given complete knowledge about the economy's data-generating process, historical economic information, and the realized shock in period t. $\mathcal{B}_{CB}$ is the central bank's inflation point projection based on an identical information set. Parameters ψ and ω denote the precision (or the reciprocal of the variance) of the distributions of δ and ν, respectively.

Agent i updates her prior of π_{t+s} with the central bank's communicated projection. Her posterior expectation is a linear combination of her private belief and the central bank's public belief,

$$\mathbb{E}_{i,t}\{\pi_{t+s} \mid \mathcal{B}_{CB}, \mathcal{B}_i\} = \frac{\psi\mathcal{B}_i + \omega\mathcal{B}_{CB}}{\psi + \omega}. \tag{4.15}$$

This Bayseian updating exercise reveals that increasing the precision of the central bank's inflation projection, ω, leads an agent to more heavily weight that projection when forming expectations. This implies that we should see more ex-ante rational forecasters in the Point treatment where the central bank's ex-ante rational projections are more precise than in the Point+Density treatment. We show in Figure 4.2 using simulations that shifting boundedly-rational subjects toward ex-ante rationality will always reduce forecaster errors and, for all but inflation targeters, reduce price volatility. Thus, we hypothesise that surrounding point forecasts with uncertainty will lead to larger forecast errors and more price volatility.

Less sophisticated agents may not fully internalize the uncertainty present in the environment. If so, then communicating uncertainty alongside the central bank's point projection can make the uncertainty salient. Depending on how the central bank communicates its projection, it may influence the perceived value of ω to boundedly-rational agents. For such agents, density projections make salient the model uncertainty surrounding the central bank's point forecast, while point-only projections may obscure the degree of uncertainty in the central bank's outlook.

This framework yields insight into how central bank communicated uncertainty influences private agents' uncertainty, measured as the conditional variance of agent i's forecast of π_{t+s},

$$Var_{i,t}(\pi_{t+s}|\, \mathcal{B}_{CB}, \mathcal{B}_i) = \frac{1}{\frac{1}{\omega} + \frac{1}{\psi}}. \tag{4.16}$$

Agent i's conditional variance about π_{t+s} is decreasing in her perception of the precision of the central bank's forecast. Thus, for a boundedly-rational agent who conflates the absence of communicated uncertainty with the absence of uncertainty surrounding the central bank's projection, providing a density forecast alongside a point forecast ought to increase individual–level uncertainty.

We also predict that central bank projections will influence the level of disagreement across forecasters. With NoComm, participants have various pieces of information at their disposal to formulate their forecasts: the DGP, historical information, and the current shock. There is no obvious focal point for coordination. By contrast, both Point and Point&Density projections provide a salient focal point projection to coordinate forecasts. Point projections provide a unique single focal point on the rational expectations equilibrium forecast and should lead to the lowest levels of disagreement. Point&Density also provides a greater focus on a one-standard-deviation range of predicted inflation values. Group treatments have a high degree of strategic complementarities; improved coordination due to the projections is also predicted to lower forecast errors.

The effects of communicating uncertainty in central bank projections on central bank credibility will depend on how credibility is measured. We can measure a participant's credibility in the forecast as their willingness to use the communicated point projection. A density forecast would 'wash out' the focal power of point projections, which undermines the ability of point projections to coordinate expectations. In this case, we would expect credibility to be higher in Point than Point&Density. Alternatively, in Point&Density, we can calculate credibility as a participant's willingness to forecast within the central bank's forecasted one-standard deviation range. With this broader definition of credibility, there is more scope for the Point&Density projection to be perceived as credible.

Projections serve two purposes to our subjects in our experiment. The projections reduce the cognitive burden associated with correctly forecasting future inflation (*information motive*). In Group treatments, aggregate expectations are the predominant driver of inflation dynamics, and a subject who aligns her expectations with aggregate (median) expectations is likely to form more accurate expectations. That is, in the Group treatments, inflation projections also provide a salient focal point to coordinate forecasts (*strategic motive*). From this perspective, we expect that projections will be more widely used and be more effective at managing and coordinating expectations in the Group treatments than Individual treatments. This logic aligns with the idea that agents react strongly to public information in environments featuring strategic complementarities, since public signals in these settings are information of the actions of others (i.e. they reduce strategic uncertainty).

On the other hand, the strategic motive to use the projections may be absent in Group treatments if participants do not believe the median forecaster(s) will incorporate the communication into their private forecasts. Participants' best response would be to incorporate their perception of aggregate expectations into their forecast. Furthermore, any increase in central bank uncertainty should make Group participants less confident that the median forecaster will adhere to the central bank's forecast. In this case, projections may be more effective at managing expectations in the Individual treatments than Group treatments.

It is not immediately apparent whether strategic considerations present in the Group treatments will lead to an increased usage of the central bank projections. However, given past experimental evidence that central bank communication can be a useful coordination device (Cornand and Heinemann, 2014; Mokhtarzadeh and Petersen, 2020), we expect the strategic motive to be sufficiently powerful. Increased credibility in the projections in the Group treatments should then lead to more ex-ante rational forecasting, less inflation volatility, and lower forecast errors.

There are some key differences between Individual and Group hypotheses, specifically in the

NoComm treatment. Without strategic concerns or an obvious focal point, Individual participants are likely to exhibit more heterogeneity in their inflation forecasts. This, in turn, can lead to higher mean inflation variability if some participants employ more extreme extrapolative forecasting heuristics in Individual than Group. At the same time, participants in Individual have full information about the aggregate forecasts influencing inflation (as it is their own forecasts) and, consequently, should form smaller forecast errors and be more confident about their own expectations than Group participants.

Furthermore, we predict that private uncertainty will be lower in the Individual than in the Group treatments. This is because Individual participants face no uncertainty about the aggregate forecast and thus the expected path of inflation. Similarly, in the absence of strategic complementarities, Individual participants are less likely to be coordinated in their forecasts and so should exhibit more disagreement.

We summarize these aggregate and individual-level predictions in the following hypotheses:

H1.a Inflation volatility$_{NoComm}$ > Inflation volatility$_{Point\&Density}$ > Inflation volatility$_{Point}$

H1.b Inflation volatility$_{Group,NoComm}$ > Inflation volatility$_{Individual,NoComm}$

H1.c Inflation volatility$_{Individual,Point,Point\&Density}$ > Inflation volatility$_{Group,Point,Point\&Density}$

H2.a Forecast errors$_{NoComm}$ > Forecast errors$_{Point\&Density}$ > Forecast errors$_{Point}$

H2.b Forecast errors$_{Group,NoComm}$ > Forecast errors$_{Individual,NoComm}$

H2.c Forecast errors$_{Individual,Point,Point\&Density}$ > Forecast errors$_{Group,Point,Point\&Density}$

H3.a Disagreement$_{NoComm}$ > Disagreement$_{Point\&Density}$ > Disagreement$_{Point}$

H3.b Disagreement$_{Individual}$ > Disagreement$_{Group}$

H4.a Uncertainty$_{NoComm}$ > Uncertainty$_{Point\&Density}$ > Uncertainty$_{Point}$

H4.b Uncertainty$_{Group}$ $>$ Uncertainty$_{Individual}$

H5.a $Credibility_{Point} > Credibility_{Point\&Density}$

H5.b $Credibility_{Group} >$ Credibility$_{Individual}$

4.3 Results

We begin by describing how point and density projections influence aggregate dynamics. We then explore how the projections influence individual forecasting behavior.

4.3.1 Aggregate results

Figure 4.3 and Figure 4.4 compare the time series of inflation for groups and individuals, respectively, across our three information treatments by sequence and repetition.[8] Time series comparing Group and Individual treatments can be found in the Online Appendix Section C. While the variability of inflation certainly differs across treatments, impressive is the contemporaneous correlation of inflation across treatments and independent groups of participants who face the same shock sequence.

We consider two measures of macroeconomic stability at the session-repetition level. First, we compute the mean deviation of inflation from the central bank's target of zero. Second, we compute the standard deviation of inflation. The mean values of both measures are presented in the first two columns of Table 4.2. Both metrics indicate that inflation variability is greatest in NoComm, followed by Point&Density, and lowest in Point. A series of Wilcoxon rank-sum tests fails to reject the null hypothesis that the distributions of these statistics are not different across treatments ($N = 6$ per treatment; $p > 0.12$ in all treatment-repetition pairwise comparisons). Our results remain qualitatively similar when we instead normalize the session-level standard deviation measures by the standard deviation of the realized shocks, which differ across sessions. Overall, we are unable to find support for Hypothesis 1a that either type of projection reduces inflation variability in Group settings.

[8] We use the terms sequence and session interchangeably.

Figure 4.3: Time series of Group treatments

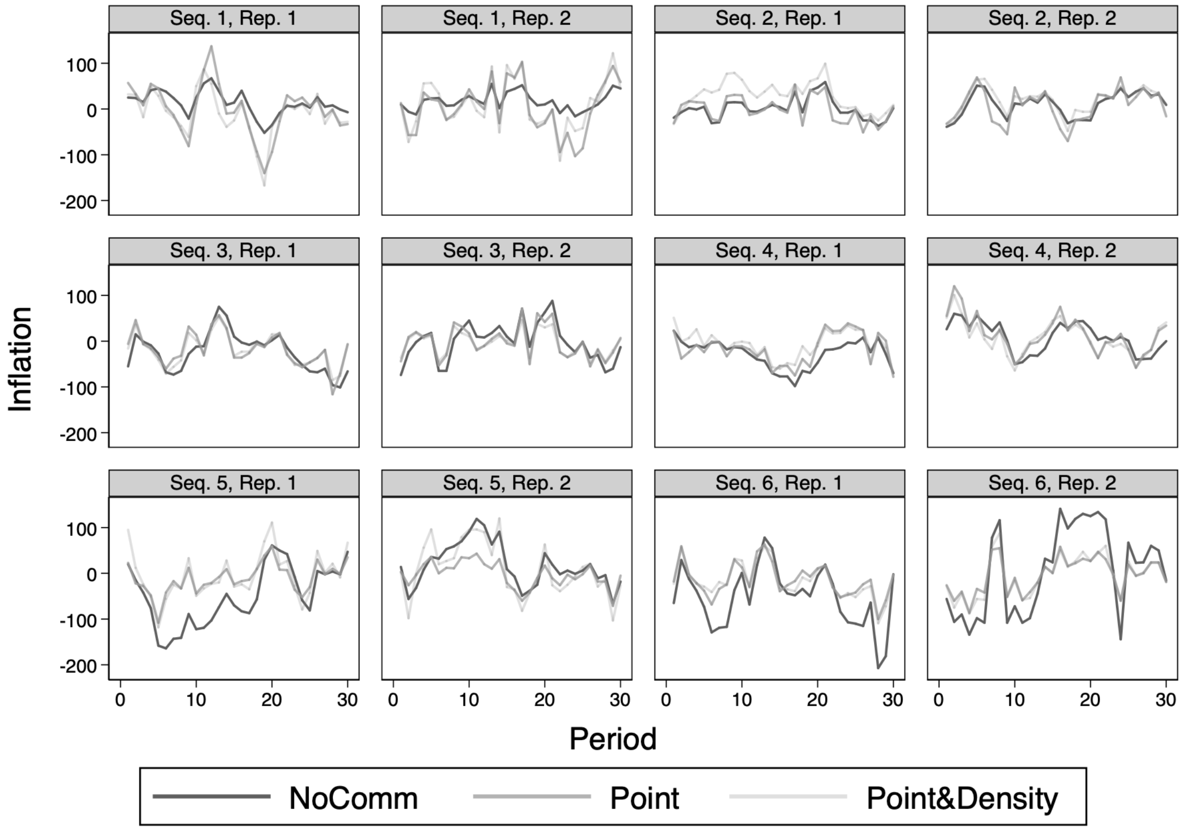

In the Individual sessions, we obtain the same ordering of treatments with NoComm exhibiting the most inflation volatility (71 bps), followed by Point&Density (62 bps) and Point (59 bps). The differences between NoComm and Point are statistically significant in Rep. 2 ($N = 39$ in NoComm, $N = 42$ in Point; $p = 0.02$ for both the raw and normalized standard deviation measures). All other treatment-repetition differences are not statistically significant ($p > 0.12$). We find minimal support for Hypothesis 1a that either type of projection reduces inflation variability in Individual settings.

We do not find significant support for Hypothesis 1b and 1c that inflation volatility is significantly different across Group and Individual treatments. While mean inflation variability is more than 50% greater in Individual treatments than Group treatments, there is considerable vari-

Figure 4.4: Time series of Individual treatments

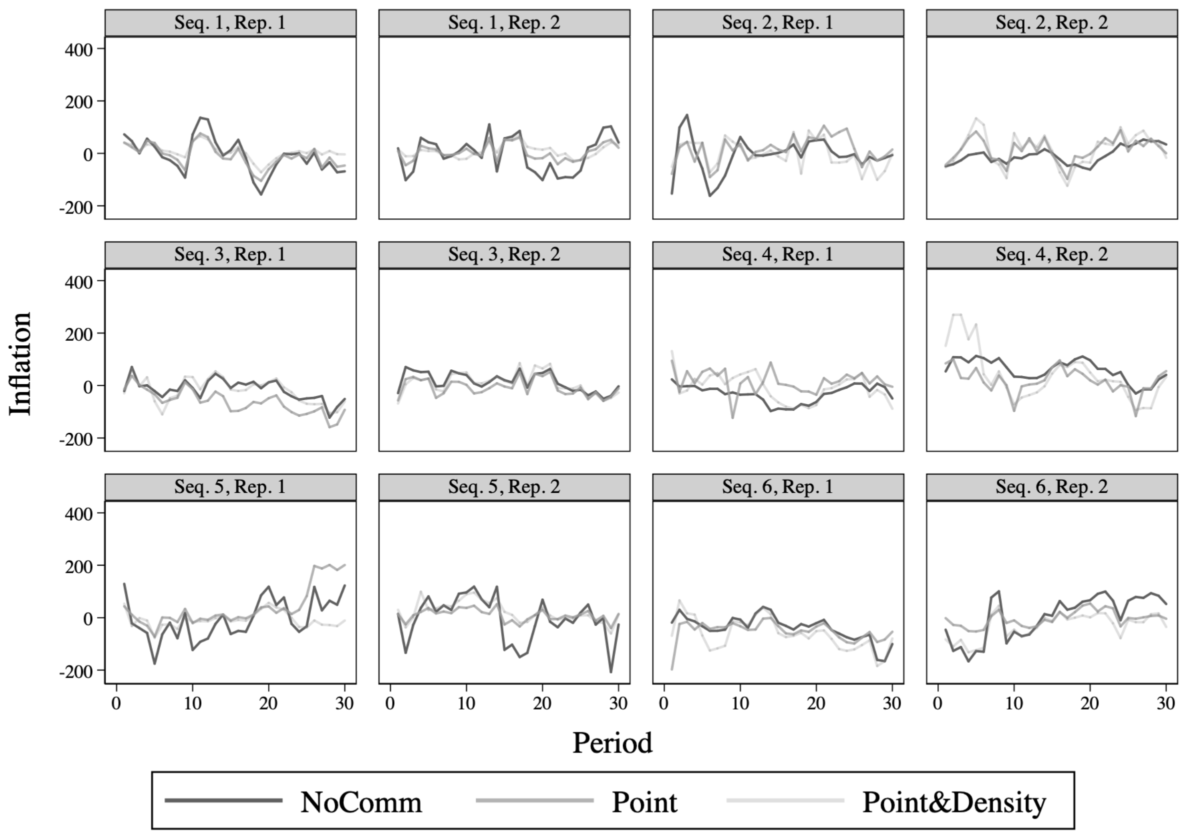

ance across Individual subjects within any given treatment. The differences between Group and Individual treatments for a given information set are not statistically significant ($N = 6$ in Group treatments and $N \geq 34$ in Individual treatments; $p > 0.17$ in all treatment-repetition comparisons).

Result 1: In Group settings, projections do not significantly improve inflation stability.

Result 2: In Individual settings, only point projections significantly reduce inflation variability for experienced participants.

4.3.2 Individual results

We now turn to our individual–level forecast data to identify participants' ability and forecasting strategies. We keep data only from those participants whose forecasts are within ± 1500 bps. This excludes five participants from each of the Point–Indiv. and Point&Density–Indiv. treatments.

Forecast Errors

Distributions of the forecast errors are presented in Figure 4.5 by repetition and coordination type, with the densities truncated at 600 for better clarity. Given the minimal differences in the distribution of forecast errors across treatments, we henceforth pool data from the two repetitions together. Forecast summary statistics of individual inflation forecasts are presented in Table 4.2. The third and fourth columns present the mean and standard deviation of absolute forecast errors of $t + 1$ and $t + 2$ inflation by treatment.

Figure 4.5: Distributions of absolute forecast errors

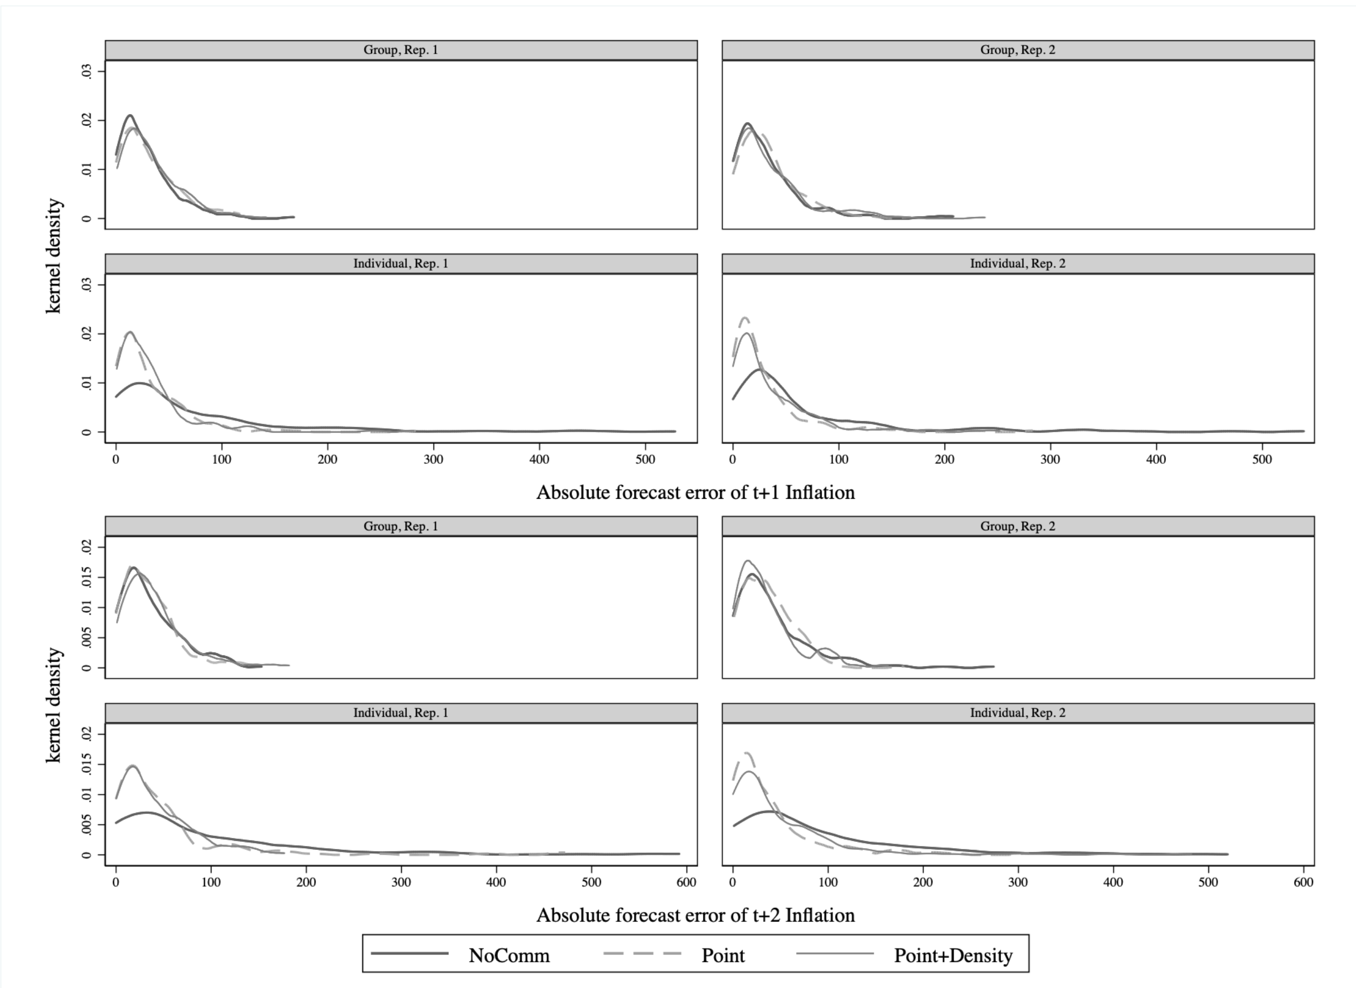

Table 4.2: Summary statistics of aggregate and individual forecast variables

Group										
CB Projection	**Deviation from Target**	**Std. Dev. Inflation**	**Abs.FE** π_{t+1}	**Abs.FE** π_{t+2}	**Disagreement** π_{t+1}	**Disagreement** π_{t+2}	**Uncertainty** π_{t+1}	**Uncertainty** π_{t+2}	**Credibility** π_{t+1}	**Credibility** π_{t+2}
NoComm	39 (37)	43 (22)	36 (56)	43 (55)	32 (46)	31 (40)	27 (37)	32 (92)	15% (0.35)	16% (0.36)
Point	32 (24)	38 (9)	31 (28)	35 (27)	17 (14)	16 (10)	17 (17)	21 (24)	41% (0.49)	41% (0.49)
Point&Density	34 (26)	40 (11)	34 (31)	38 (35)	21 (17)	21 (20)	30 (29)	35 (32)	34% (0.47) 99%† (0.08)	37% (0.48) 99%† (0.05)
Individual										
CB Projection	**Deviation from Target**	**Std. Dev. Inflation**	**Abs.FE** π_{t+1}	**Abs.FE** π_{t+2}	**Disagreement** π_{t+1}	**Disagreement** π_{t+2}	**Uncertainty** π_{t+1}	**Uncertainty** π_{t+2}	**Credibility** π_{t+1}	**Credibility** π_{t+2}
NoComm	66 (79)	117 (89)	43 (68)	57 (80)	74 (44)	73 (45)	23 (45)	26 (35)	11% (0.11)	11% (0.32)
Point	51 (93)	53 (41)	37 (68)	43 (82)	66 (70)	62 (68)	19 (29)	21 (29)	36% (0.48)	35% (0.48)
Point&Density	58 (73)	65 (60)	36 (51)	47 (62)	65 (45)	61 (45)	24 (28)	26 (30)	23% (0.42) 86%† (0.35)	24% (0.43) 91%† (0.28)

This table presents means and standard deviation for each variable by treatment. Units are given in basis points, except for Credibility which is the percentage of participants who forecast the central bank's projected value within 5 bps. † denotes the percentage of forecasts that fall within the central bank's projected range in the PointDensity treatment.

We find mixed support for Hypothesis 2a that projections reduce forecast errors, with Point projections more effective than Point&Density projections at improving forecast accuracy. Consistent across both groups and individuals, as well as one- and two-period-ahead forecasts, we find that absolute forecast errors are largest in the NoComm, followed by Point&Density, and lowest in the Point treatments.

Table 4.3: Absolute forecast errors

Panel A: Information comparisons

	Group				Individual			
	$t+1$		$t+2$		$t+1$		$t+2$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Point	-5.057		-7.829**		-6.460		-13.082*	
	(3.37)		(3.32)		(5.71)		(6.99)	
Point&Density	-2.155	2.902*	-4.795	3.034*	-7.376	-0.916	-9.759	3.322
	(3.51)	(1.52)	(3.47)	(1.72)	(4.94)	(4.93)	(6.61)	(6.04)
α	35.901***	30.844***	42.690***	34.862***	43.161***	36.701***	56.678***	43.596***
	(3.26)	(0.82)	(3.17)	(1.00)	(4.05)	(4.03)	(5.30)	(4.56)
N	7306	4872	7054	4704	6604	4343	6377	4194
χ^2	5.237	3.629	7.437	3.094	2.336	0.0346	3.684	0.303

Panel B: Group vs. Individual comparisons

	NoComm		Point		Point&Density	
	$t+1$	$t+2$	$t+1$	$t+2$	$t+1$	$t+2$
	(1)	(2)	(3)	(4)	(5)	(6)
Individual	7.260	13.986**	5.858	8.735*	2.040	9.023**
	(5.21)	(6.18)	(4.11)	(4.67)	(3.11)	(4.20)
α	35.901***	42.691***	30.844***	34.862***	33.746***	37.896***
	(3.27)	(3.17)	(0.82)	(1.00)	(1.29)	(1.41)
N	4695	4533	4720	4558	4495	4340
χ^2	1.946	5.120	2.028	3.500	0.430	4.623

This table presents results from a series of random effects panel regressions. Units are given in basis points. The dependent variables are the absolute one- and two-period-ahead absolute forecast errors of inflation. Point, Point&Density, and Individual are treatment-specific dummy variables. α denotes the estimated constant. Robust standard errors are given in parentheses. *$p < 0.10$, **$p < 0.05$, and ***$p < 0.01$.

To evaluate whether the differences are statistically significant, we conduct a series of random effects panel regressions where we regress absolute forecast errors on treatment-specific dummy variables. Table 4.3 Panel A presents the results for Group treatments in columns (1) to (4) and Individual treatments in columns (5) to (8). Odd columns compare the two projection treatments to the NoComm treatment (denoted by α). The even columns compare forecast errors in Point&Density to Point. While forecast errors do decline with projections, the effect is not statistically significant in the Group treatments for one-period-ahead forecasts. Two-period-ahead inflation forecast errors are significantly lower when a Point projection is communicated. Columns (2) and (4) show that adding a density forecast to an existing point forecast can lead to a small but statistically significant increase in both one- and two-period-ahead forecast errors. In the Individual treatments, point projections significantly decrease two-period-ahead forecast errors by roughly 14 bps. Overall, however, the projections do not have a consistent effect on one-period-ahead forecasts.

Table 4.3 Panel B presents the estimated effects of eliminating strategic motives on absolute forecast errors, by treatment. $Individual$ is a dummy variable that takes the value of one if participants are in the Individual treatment, with the Group treatment taken as the baseline. Consistently, two-period-ahead forecast errors are more extreme in Individual treatments than in Group treatments. This difference is statistically significant at the 5% level in the NoComm and Point&Density treatments, and 10% level in the Point treatment. On average, one-period-ahead forecast errors are also larger in the Individual treatments, but the effect is not statistically significant. Thus, we reject Hypothesis 2b that forecast errors in NoComm-Indiv. are lower than in No-Comm Group, and find some support for Hypothesis 2c that errors are smaller in Group treatments with central bank projections.

Result 3: Point projections significantly reduce $t + 2$ ahead forecast errors, and are significantly more effective than Point&Density projections in Group treatments.

Result 4: Participants in Individual treatments form significantly larger forecast errors about $t + 2$ inflation than their Group counterparts.

Results 3 and 4 coincide with our hypotheses. The fact that projections are more effective at reducing forecast errors for $t + 2$ than $t + 1$ is likely due to the additional cognitive complexity associated with forecasting further into the future. Regarding Result 4, we note that there are more outlier forecasts, fewer ex-ante rational subjects and considerably more trend-chasing heuristics in Individual than Group treatments. Thus, it is not surprising that the forecast errors are larger in Individual treatments.

Disagreement

We next consider how forecast disagreement is affected by the communication of projections. We measure forecast disagreement at the session-period level as the standard deviation of inflation forecasts across subjects. Mean and standard deviations of forecast disagreement are presented in the third and fourth columns of Table 4.2. Table 4.4 provides estimates of the treatment differences in disagreement.

Table 4.4: Disagreement in inflation forecasts

Panel A: Information comparisons

	Group				Individual			
	$t+1$		$t+2$		$t+1$		$t+2$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Point	-14.937**		-15.685**		-8.346		-11.102	
	(7.16)		(6.31)		(11.41)		(9.96)	
Point&Density	-10.616	4.321	-10.038	5.647*	-9.342	-0.997	-11.406	-0.304
	(7.29)	(2.83)	(6.63)	(3.13)	(10.80)	(12.50)	(10.45)	(11.37)
α	31.634***	16.697***	31.259***	15.574***	74.220***	65.875***	72.889***	61.787***
	(6.95)	(1.76)	(6.08)	(1.68)	(6.80)	(9.22)	(6.36)	(7.72)
N	1080	720	1080	720	1080	720	1080	720
χ^2	5.877	2.335	8.368	3.245	0.941	0.00635	1.753	0.000716

Panel B: Group vs. Individual comparisons

	NoComm		Point		Point&Density	
	$t+1$	$t+2$	$t+1$	$t+2$	$t+1$	$t+2$
	(1)	(2)	(3)	(4)	(5)	(6)
Individual	42.586***	41.630***	49.178***	46.213***	43.860***	40.262***
	(9.79)	(8.86)	(9.38)	(7.90)	(8.73)	(8.76)
α	31.634***	31.259***	16.697***	15.574***	21.017***	21.221***
	(7.00)	(6.13)	(1.76)	(1.68)	(2.21)	(2.65)
N	720	720	720	720	720	720
χ^2	18.91	22.07	27.46	34.22	25.22	21.13

This table presents results from a series of random effects panel regressions. Units are given in basis points. The dependent variables are the per-period standard deviations of one- and two-period-ahead forecasts of inflation, computed at the session level. Point, Point&Density, and Individual are treatment-specific dummy variables. α denotes the estimated constant. Robust standard errors are given in parentheses. *$p < 0.10$, **$p < 0.05$, and ***$p < 0.01$.

We find mixed support for Hypothesis 3a that Point and Point&Density projections reduce disagreement. Within Group treatments, we find that Point projections significantly reduce both one- and two-period-ahead disagreement ($p < 0.05$ in both cases). We also find that the additional inclusion of densities around a point projection leads to a small but significant increase in disagreement in two-period-ahead disagreement. Point&Density projections, when compared to NoComm, do not significantly reduce disagreement. Within Individual treatments, the two types of projections reduce disagreement across subjects by roughly ten bps, but the differences across Communication treatments are not statistically significant at the 10% level.

Disagreement across subjects is significantly higher in the Individual treatments when no strategic coordination motive is present ($p < 0.001$ in all communication treatments). Disagreement falls by more than 50% in NoComm Groups, by 74% in Point Groups, and by roughly 65% in Point+Density Groups. Thus, we fail to reject Hypothesis 3b that disagreement is higher in Individual treatments.

Result 5: Point projections significantly reduce disagreement about future inflation, but Point&Density projections are not consistently effective.

Result 6: Disagreement is significantly lower in the Group treatment than in the Individual treatment, i.e. when participants have a strategic motive to coordinate their forecasts.

Uncertainty

Subjects provided predictions of their forecast errors, which we take as a measure of uncertainty. Mean and standard deviations of expected forecast errors are presented in the final two columns of Table 4.2. Table 4.5 provides estimates of the treatment differences in participants' conveyed uncertainty.

Table 4.5: Uncertainty in inflation forecasts

Panel A: Information comparisons

	Group				Individual			
	$t+1$		$t+2$		$t+1$		$t+2$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Point	-9.377***		-11.790**		-3.858		-5.129	
	(3.53)		(4.70)		(3.26)		(3.61)	
Point&Density	3.650	13.027***	1.938	13.728***	0.305	4.163	0.089	5.218
	(4.06)	(2.92)	(4.98)	(3.52)	(3.47)	(3.09)	(3.82)	(3.39)
α	26.703***	17.326***	32.894***	21.105***	23.339***	19.482***	26.047***	20.918***
	(3.20)	(1.49)	(4.15)	(2.19)	(2.56)	(2.02)	(2.85)	(2.22)
N	7559	5040	7559	5040	6840	4500	6840	4500
χ^2	22.96	19.92	17.30	15.20	2.301	1.811	3.142	2.376

Panel B: Group vs. Individual comparisons

	NoComm		Point		Point&Density	
	$t+1$	$t+2$	$t+1$	$t+2$	$t+1$	$t+2$
	(1)	(2)	(3)	(4)	(5)	(6)
Individual	27.817**	19.853*	1.017	0.656	-10.045*	-13.212*
	(11.53)	(10.35)	(8.68)	(9.47)	(6.08)	(7.57)
α	18.392***	23.639***	20.903***	22.889***	29.136***	34.122***
	(1.81)	(3.47)	(4.90)	(5.68)	(4.92)	(6.39)
N	720	720	720	720	720	720
χ^2	5.822	3.681	0.0137	0.00479	2.732	3.042

This table presents results from a series of random effects panel regressions. Units are given in basis points. The dependent variables are the participants' expected errors in their one- and two-period-ahead forecasts of inflation. Point, Point&Density, and Individual are treatment-specific dummy variables. α denotes the estimated constant. Robust standard errors are given in parentheses. $*p < 0.10$, $**p < 0.05$, and $***p < 0.01$.

Immediately striking is the high level of confidence participants convey alongside their point forecasts. The average uncertainty is approximately 30 bps in NoComm, 20 bps in Point, and 33 bps in Point&Density in Group settings and slightly lower in Individual. A rational agent asked to convey a one-standard deviation forecast would predict an expected error of 113 bps. Our participants are much more confident in their forecasts, conveying only one–quarter of one–standard–deviation of a rational level of uncertainty.[9]

We predicted in Hypothesis 4a that uncertainty would be the highest in the NoComm treatment, followed by Point& Density, and lowest in the Point treatment. We find significant support for this hypothesis in both the Group and Individual treatments. Within Group treatments, one- and two-period-ahead uncertainty decreases by approximately ten bps when the central bank communicates a Point projection. This effect is significant at the 1% (5%) level for one- (two-) period ahead forecasts. Communicating an auxiliary density around the point projection significantly increases both forecast uncertainties by approximately 14 bps. This effect is significant at the 1% level. We obtain qualitatively similar effects from projections in the Individual treatments, though the effects are smaller and not statistically significant.

In Hypothesis 4b we predicted that introducing strategic considerations would increase participants' uncertainty about future inflation. We find mixed evidence to support this. Only in the Point&Density treatment are subjects significantly more unsure about their personal forecasts when dealing with other participants. In Point and, especially, NoComm, strategic coordination leads to less uncertainty about future inflation.

Result 7: Point projections significantly reduce individual-level uncertainty about forecasts of future inflation in Group treatments, but Point&Density projections are not consistently effective.

Result 8: Individuals exhibit less uncertainty about their private forecasts than Groups

[9] Pfajfar and Žakelj (2016) also observe a high level of overconfidence in related LtF New Keynesian experiments. Uncertainty declines as the central bank pursues a more aggressive reaction to deviations of inflation from target and induces more stable inflation dynamics.

when presented with Point&Density projections.

Credibility

Credibility is an important concern for central banks who communicate their projections to the public. We denote a participant as perceiving the central bank's projection as credible if she uses its projected point forecast to formulate her private expectations. Given the potential for rounding errors, we assume a participant uses the projected value if their forecast is within five basis points of the projection. Table 4.6 provides estimates of the treatment differences in participants' credibility in the central bank's projections.

Without any communication, roughly 15% (11%) of one- and two-period-ahead forecasts in NoComm-Group (NoComm-Indiv.) are within five basis points of the rational expectations equilibrium forecast. Point projections are used by 41% (36%) of Group (Indiv.) subjects to formulate their one- and two-period-ahead forecasts. Communicating a density decreases the usage of the point projection. Credibility in the actual point prediction decreases to 34% (23%) for one-period-ahead forecasts and 37% (24%) for two-period-ahead forecasts in the Group (Indiv.) treatments.

We can alternatively consider credibility in the Point&Density treatment to include any forecast in the central bank's communicated density forecast. Credibility according to this definition is 99% for both one- and two-period ahead forecasts in the Group treatments, and between 86% and 91% in the Individual treatments. For reference, NoComm and Point also exhibit nearly identical levels of credibility for both variables subjects forecast. This high degree of similarly across information treatments indicates that the density projection is not improving credibility, and, if anything, is reducing credibility in the central bank's *point* projections.

Table 4.6: Credibility of central bank projections

Panel A: Information comparisons

	Group				Individual			
	$t+1$		$t+2$		$t+1$		$t+2$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Point	0.267***		0.156***		0.270***		0.146***	
	(0.03)		(0.02)		(0.03)		(0.02)	
Point&Density	0.196***	-0.071*	0.082***	-0.074***	0.140***	-0.130***	0.059***	-0.087***
	(0.03)	(0.04)	(0.02)	(0.02)	(0.03)	(0.05)	(0.02)	(0.03)
α	0.147***	0.413***	0.152***	0.308***	0.109***	0.378***	0.114***	0.260***
	(0.01)	(0.03)	(0.01)	(0.02)	(0.01)	(0.03)	(0.01)	(0.02)
N	7559	5040	7559	5040	6840	4500	6840	4500
χ^2	98.30	2.812	62.74	10.05	72.07	7.668	42.92	10.13

Panel B: Group vs. Individual comparisons

	NoComm		Point		Point&Density	
	$t+1$	$t+2$	$t+1$	$t+2$	$t+1$	$t+2$
	(1)	(2)	(3)	(4)	(5)	(6)
Individual	-0.038**	-0.038**	-0.035	-0.048*	-0.094**	-0.061**
	(0.01)	(0.02)	(0.04)	(0.03)	(0.05)	(0.02)
α	0.147***	0.152***	0.413***	0.308***	0.342***	0.234***
	(0.01)	(0.01)	(0.03)	(0.02)	(0.03)	(0.01)
N	4859	4859	4890	4890	4650	4650
χ^2	6.549	6.196	0.627	3.190	4.325	6.614

This table presents results from a series of random effects panel regressions. The dependent variables are dummy variables that take the value of one if one- and two-period-ahead forecasts of inflation are less than five basis points from the central bank's point projection. Point, Point&Density, and Individual are treatment-specific dummy variables. α denotes the estimated constant. Robust standard errors are given in parentheses. *$p < 0.10$, **$p < 0.05$, and ***$p < 0.01$.

For both Group and Individual treatments, communicating either a Point or Point&Density projection significantly increases the share of participants forecasting the REE solution for $t + 1$ inflation. Consistent with Hypothesis 5a, credibility in the central bank's point projection is significantly lower, however, when the projection includes a density forecast for both $t + 1$ and $t + 2$ forecasts in Group and Individual treatments.

We also observe a small but significant increase in credibility in the projections when participants face strategic considerations. The effect is roughly four percentage points in the Point treatments, and between 6-9 percentage points in the Point&Density treatments. Thus, we find support for Hypothesis 5b that Group settings improve credibility in projections.

Result 9: Credibility in the central bank's point forecast is significantly lower when the central bank communicates a density around its point projection.

Result 10: Credibility in projections is higher when participants interact in Group treatments.

Heuristics

Finally, we consider how the projections and strategic considerations alter the heuristics subjects use to formulate their forecasts. This exercise not only provides insight into whether projections have the intended impact on expectations, but also highlights which types of heuristics become more or less prevalent in the presence of central bank communication. Table 4.7 presents the six general classes of heuristics we consider. The heuristics have been previously identified by theory and experiments as describing forecasters' expectation formation process.

Following Mokhtarzadeh and Petersen (2020), we classify each participant into one of six heuristics that most closely matches their own submitted expectations. Specifically, we identify the heuristic that produces the lowest absolute mean-squared error among all competing models. For the Constant Gain and Trend Chasing heuristics, we consider a range of parameters $\gamma, \tau \in [0.1, 1.5]$

Table 4.7: Forecasting heuristics

Model	Heuristic Name	Model
M1	Ex-ante rational	$E_{i,t}\pi_{t+1} = 0.08r^n_{t-1} + 0.14\epsilon_t$
		$E_{i,t}\pi_{t+2} = 0.05r^n_{t-1} + 0.08\epsilon_t$
M2	Adaptive(1)	$E_{i,t}\pi_{t+1} = 0.09r_{t-1} + 0.88\pi_{t-1} + 0.17\epsilon_t$
		$E_{i,t}\pi_{t+2} = 0.10r_{t-1} + 0.84\pi_{t-1} + 0.18\epsilon_t$
M3	Target	$E_{i,t}\pi_{t+1} = 0$
		$E_{i,t}\pi_{t+2} = 0$
M4	Naive	$E_{i,t}\pi_{t+1} = \pi_{t-1}$
		$E_{i,t}\pi_{t+2} = \pi_{t-1}$
M5	Constant Gain	$E_{i,t}\pi_{t+1} = E_{i,t-2}\pi_{t-1} - \gamma(E_{i,t-2}\pi_{t-1} - \pi_{t-1})$
		$E_{i,t}\pi_{t+2} = E_{i,t-3}\pi_{t-1} - \gamma(E_{i,t-2}\pi_{t-1} - \pi_{t-1})$
M6	Trend Chasing	$E_{i,t}\pi_{t+1} = \pi_{t-1} + \tau(\pi_{t-1} - \pi_{t-2})$
		$E_{i,t}\pi_{t+2} = \pi_{t-1} + 2\tau(\pi_{t-1} - \pi_{t-2})$

Models of expectations as functions of exogenous or historical data. γ and $\tau \in [0.1, 1.5]$ in increments of 0.1.

with 0.1 increments. The distribution of $t + 1$ inflation forecasting heuristics are presented in Figure 4.6 by treatment.

There are many interesting results to be taken away from these analyses. Without any auxiliary communication, between 10 and 20% of participants in both Group and Individual treatments formulate ex-ante rational expectations. Importantly, after controlling for experience, there is little difference in the prevalence of rational agents in strategic and individual environments. This is noteworthy as one might assume that participants' potential irrationality in NoComm may be due to the perceived irrationality of their counterparts. Rather, it is in the NoComm-Individual treatment that we observe a greater frequency of highly irrational heuristics such as Trend Chasing.

Communicating a Point projection is very effective at guiding participants to forecasting the REE solution. Roughly 80% of Group participants and 48% of Individual participants behave as if they were forming ex-ante rational expectations when they receive Point projections. In Point-Individual, the inflation projection effectively nudges subjects away from using Adaptive and Trend-Chasing heuristics toward both Rational and Constant Gain. The point projection is noticeably less effective in the Individual treatment, likely because Individual participants who do not initially utilize the projection to formulate their forecast observe dynamics that look different

Figure 4.6: Distribution of forecasting heuristics for $t+1$ inflation, by treatment

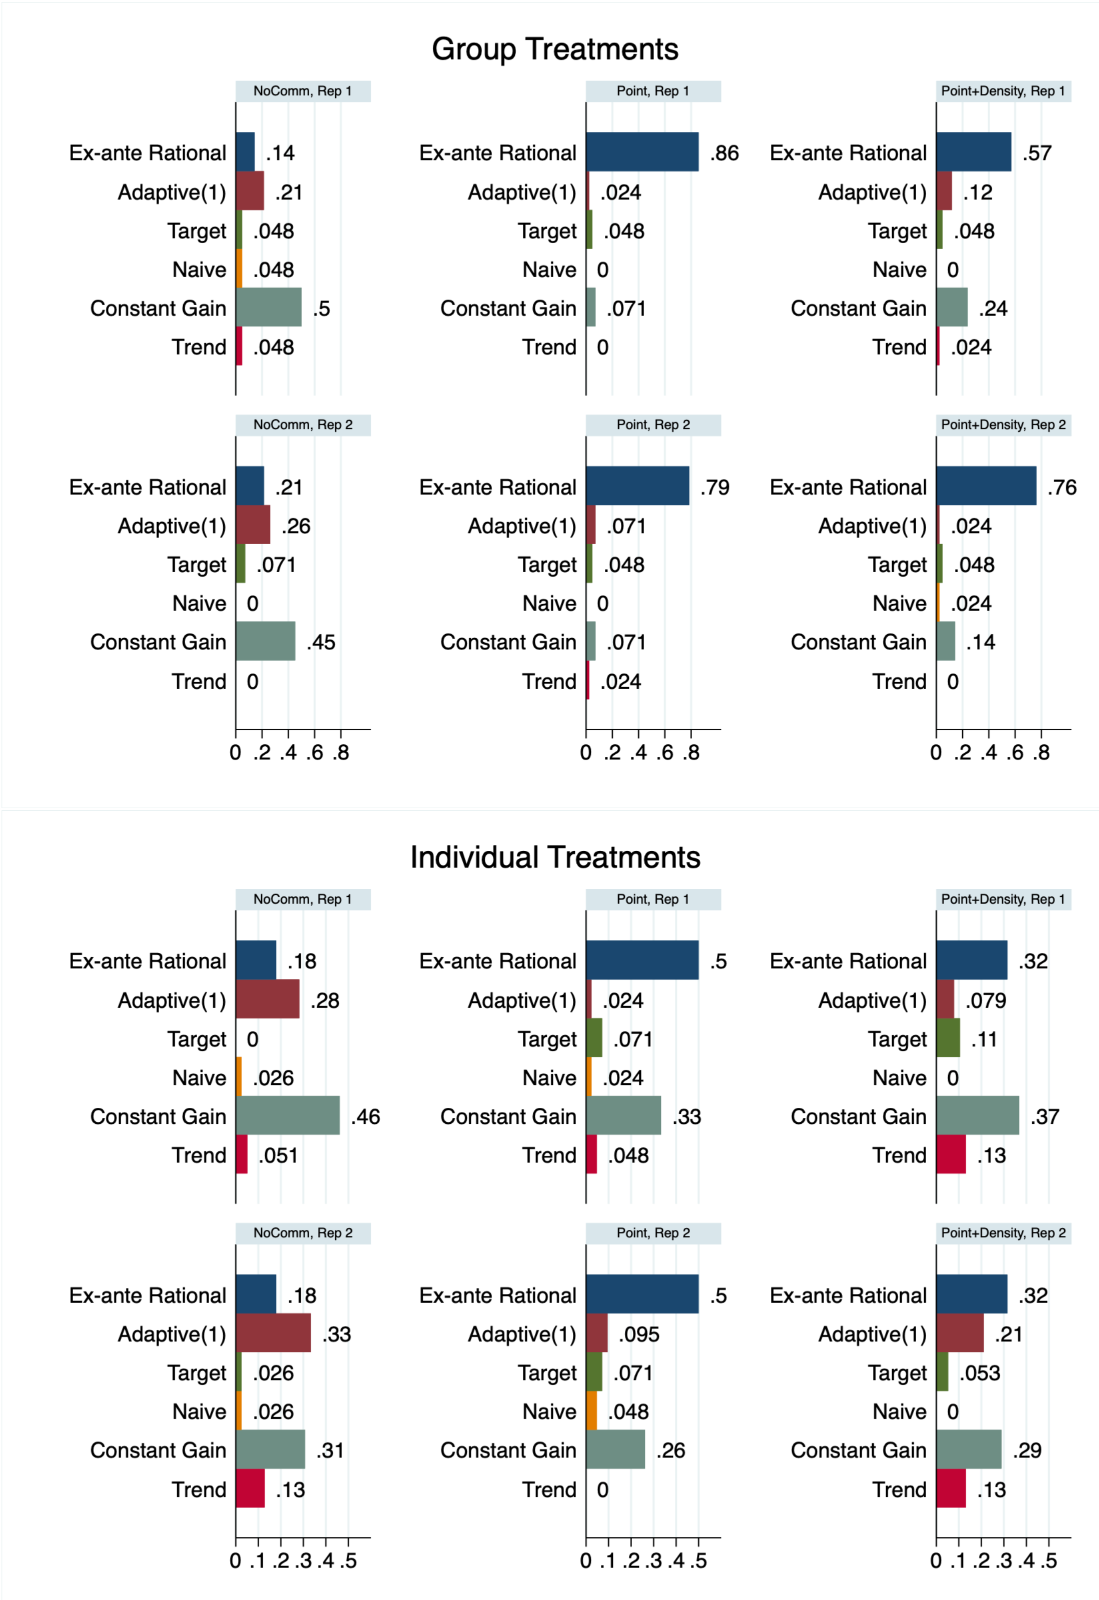

from the projected values. They subsequently lose credibility in the projections.

Density projections also increase the proportion of subjects who forecast as if they were Ex-ante Rational and reduces the proportion of Adaptive forecasters. However, for inexperienced Group participants and both inexperienced and experienced Individual participants, the inclusion of the density projection mutes the effects of the point projection. In addition to the previously noted heterogeneity in forecasts, we also observe considerably greater heterogeneity in heuristics in Point&Density compared to Point.

Between 76 and 87% of participants use the same heuristic to forecast one- and two-period-ahead inflation, without much difference across treatments. For those that exhibit differences, a few consistencies emerge. Adaptive forecasters of $t+1$ inflation tend to be Rational for their $t+2$ forecasts in projection treatments. Rational forecasters of $t+1$ inflation tend to be primarily split between Target and Trend-Chasing for their subsequent forecast. Finally, those that forecast $t+1$ inflation with a Trend Chasing heuristic use predominantly a Rational heuristic to forecast their subsequent forecast, in treatments with projections.

Result 11: Ex-ante rational projections reduce the prevalence of backward-looking forecasting heuristics and encourage more rational forecasting.

Result 12: Point projections are more effective at guiding expectations to the REE than Point&Density projections.

Result 13: The majority of participants use the same heuristics to formulate both their one- and two-period-ahead forecasts.

4.4 Conclusion

As more central banks publish forecasts about their outlook, they face the dilemma as to whether to communicate their own uncertainty. To the best of our knowledge, there has been no work evaluating the impact of publishing density forecasts in addition to point projections on market expectations.

Our work aims to fill this gap by providing original evidence on the effects of communicating uncertainty on expectation formation. First, we study the introduction of a measure of a central bank's forecast uncertainty into central bank projections (i.e. the publication of density rather than point forecasts). Our interest is in how this affects aggregate dynamics and how forecasters incorporate information in the first and second central moments into their own expectations and perceptions of future uncertainty. Second, this paper studies behavior in individual-choice and coordination settings to understand the extent to which strategic concerns influence how agents use information when forming expectations.

We find that both point and density projections significantly improve forecast accuracy and decrease cross-sectional disagreement relative to an environment with no auxiliary central bank communication. This is consistent with empirical evidence that central bank projections coordinate expectations and reduce forecast errors. Furthermore, projections increase the proportion of participants who form ex-ante rational expectations. We provide new evidence showing that a large majority of participants use the same heuristics to formulate both their one- and two-period-ahead forecasts. However, roughly 20% of participants employ different heuristics when forecasting at different horizons. These subjects tend to use more irrational heuristics for their further ahead forecasts. However, projections nudge more distant forecasts toward the ex-ante rational prediction.

These results are in line with Mokhtarzadeh and Petersen (2020) who find inflation projections works effectively to guide expectations in the absence of a zero lower bound (ZLB). Inflation expectations are not as well managed by rationally–constructed inflation projections in the presence of ZLB. See Arifovic and Petersen (2017), Ahrens et al. (2018), and Kryvtsov and Petersen (2020) for examples of relatively poorer management of inflation expectations through projections at ZLB. In particular, more simplistic communications are more effective to guide expectations. Kryvtsov and Petersen note that simple, relatable information about past interest rates work more effectively to manage expectations than forward-looking projections about policy rates and forward guidance. Ahrens et al. find that gradual adjustment of inflation projections by human central bankers can more effectively build up credibility and manage expectations at the ZLB. Arifovic and Petersen find that qualitative rather than quantitative communication can work somewhat better to reduce pessimism.

Communicating an additional density forecast around a point projection mutes the positive effects of publishing point projections. Compared to point projections, density projections significantly increase forecast errors and disagreement. The central bank transmits their uncertainty to forecasters, leading to higher levels of private forecast uncertainty. Moreover, fewer subjects form ex-ante rational expectations at both forecasting horizons.

Credibility in the central bank's projections significantly decreases when participants are presented with a less precise projection. This result is in line with Baeriswyl and Cornand (2016) who show in a Keynesian beauty contest environment that subjects place more weight on public signals the more precise is the signal. Their results are more nuanced. Subjects overreact relative to theoretical predictions when the public signal is imprecise, and under-react when it is more precise.

A notable finding in our experiment is that inflation volatility and the heterogeneity in forecasts and heuristics increases when participants have more market power in the Individual treatments. Other LTF experiments have also explored the effects of individual subjects' market size on system stability. Kopányi et al. (2019) show in an asset market LTF that increasing the market power of more accurate forecasters can lead to greater instability. Bao et al. (2020), on the other hand, find that asset price bubbles grow even faster with larger group sizes (where individuals have less market power) and participants are more likely to coordinate on trend-chasing strategies. The differences in our experimental findings and Bao et al. likely are driven by the relatively greater negative feedback present in our data-generating process, that encourage coordination on more stable heuristics (see Heemeijer et al. 2009 for evidence on the effects of positive and negative feedback in LTF experiments).

We show that central bank projections both provide valuable information to reduce forecasters' confusion and can alleviate some strategic uncertainty. Our results line up with Akiyama et al. (2017) who show that mispricing commonly observed in experimental asset markets is driven at least in part by strategic uncertainty. In fact, they show that strategic uncertainty explains at least as much of median initial forecast deviation from the fundamental value as does confusion. Our results are consistent with the broader literature that shows individuals will anchor on simple and salient information when facing cognitive overload (Deck and Jahedi, 2015, Kryvtsov and Petersen, 2020).

Much of the macroeconomic literature uses a notion of rational expectations that focuses on an agents' point forecast of some future event. There is scope in learning-to-forecast experiments to better understand participants' subjective uncertainty, how it relates to rationality, and how the relationship between rationality and uncertainty can be influenced by monetary policy and central bank communication. We find that participants exhibit an unusually high level of confidence in their own forecasts, and this confidence can be better strengthened through more precise information. An important avenue of future study is whether participants would act on expectations given their level of confidence.

4.5 References

1. Ahrens, S., Lustenhouwer, J., & Tettamanzi, M. (2018). The Stabilizing Role of Forward Guidance: A Macro Experiment. BERG Working Paper 137.

2. Akiyama, E., Hanaki, N., & Ishikawa, R. (2017). It is Not Just Confusion! Strategic Uncertainty in An Experimental Asset Market. Economic Journal, 127(605), 563-580.

3. Arifovic, J., & Petersen, L. (2017). Stabilizing expectations at the zero lower bound: Experimental evidence. Journal of Economic Dynamics and Control, 82, 21-43.

4. Assenza, T., Heemeijer, P., Hommes, C. H., & Massaro, D. (2013). Individual expectations and aggregate macro behavior.

5. Baeriswyl, R., & Cornand, C. (2016). The predominant role of signal precision in experimental beauty contests. The BE Journal of Theoretical Economics, 16(1), 267-301.

6. Bao, T., Duffy, J., & Hommes, C. (2013). Learning, forecasting and optimizing: An experimental study. European Economic Review, 61, 186-204.

7. Bao, T., Hennequin, M., Hommes, C., & Massaro, D. (2020). Coordination on bubbles in large-group asset pricing experiments. Journal of Economic Dynamics and Control, 110, 103702.

8. Bauer, M. D. (2012). Monetary policy and interest rate uncertainty. FRBSF Economic Letter, 38, 1-5.

9. Blinder, A. S., Ehrmann, M., Fratzscher, M., De Haan, J., & Jansen, D. J. (2008). Central bank communication and monetary policy: A survey of theory and evidence. Journal of Economic Literature, 46(4), 910-45.

10. Brubakk, L., Ter Ellen, S., & Xu, H. (2017). Forward guidance through interest rate projections: does it work? Norges Bank Working Paper 6/2017.

11. Coibion, O., & Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. American Economic Review, 105(8), 2644-78.

12. Cornand, C., & Heinemann, F. (2008). Optimal degree of public information dissemination. The Economic Journal, 118(528), 718-742.

13. Cornand, C., & Hubert, P. (2019). On the external validity of experimental inflation forecasts: A comparison with five categories of field expectations. Journal of Economic Dynamics and Control, 103746.

14. Cornand, C., & M'baye, C. K. (2018a). Does inflation targeting matter? An experimental investigation. Macroeconomic Dynamics, 22(2), 362-401.

15. Cornand C. & M'baye C. K. (2018b), "Band or Point Inflation Targeting? An Experimental Approach", Journal of Economic Interaction and Coordination, 13(2) : 283-309.

16. Deck, C., & Jahedi, S. (2015). The effect of cognitive load on economic decision making: A survey and new experiments. European Economic Review, 78, 97-119.

17. Geraats, P. M. (2002). Central bank transparency. The Economic Journal, 112(483), F532-F565.

18. Gosselin, P., Lotz, A., & Wyplosz, C. (2008). The Expected Interest Rate Path: Alignment of Expectations vs. Creative Opacity. International Journal of Central Banking, 4(3), 145-185.

19. Greenspan, A. (2004). Risk and uncertainty in monetary policy. American Economic Review, 94(2), 33-40.

20. Heemeijer, P., Hommes, C., Sonnemans, J., Tuinstra, J. (2009). Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. Journal of Economic dynamics and control, 33(5), 1052-1072.

21. Hommes, C., Massaro, D., & Weber, M. (2019a). Monetary policy under behavioral expectations: Theory and experiment. European Economic Review, 118, 193-212.

22. Hommes, C., Massaro, D., & Salle, I. (2019b). Monetary and fiscal policy design at the zero lower bound: Evidence from the lab. Economic Inquiry, 57(2), 1120-1140.

23. Hubert, P. (2014). FOMC forecasts as a focal point for private expectations. Journal of Money, Credit and Banking, 46(7), 1381-1420.

24. Hubert, P. (2015). Do central bank forecasts influence private agents? Forecasting performance versus signals. Journal of Money, Credit and Banking, 47(4), 771-789.

25. Husted, L., Rogers, J., & Sun, B. (2018). Uncertainty, currency excess returns, and risk reversals. Journal of International Money and Finance, 88, 228-241.

26. Husted, L., Rogers, J., & Sun, B. (2019). Monetary policy uncertainty. Journal of Monetary Economics, 115, 20-36.

27. Jain, M., & Sutherland, C. S. (2018). How do central bank projections and forward guidance influence private-sector forecasts? (No. 2018-2). Bank of Canada Staff Working Paper.

28. Kool, C. J., & Thornton, D. L. (2015). How effective is central bank forward guidance?

29. Kopányi, D., Rabanal, J. P., Rud, O. A., & Tuinstra, J. (2019). Can competition between forecasters stabilize asset prices in learning to forecast experiments? Journal of Economic Dynamics and Control, 109, 103770.

30. Kryvtsov, O., & Petersen, L. (2013). Expectations and monetary policy: experimental evidence (No. 2013-44). Bank of Canada.

31. Kryvtsov, O., & Petersen, L. (2020). Central bank communication that works: Lessons from lab experiments. Journal of Monetary Economics. Forthcoming.

32. Malmendier, U., & Nagel, S. (2016). Learning from inflation experiences. The Quarterly Journal of Economics, 131(1), 53-87.

33. Mokhtarzadeh, F., & Petersen, L. (2020). Coordinating expectations through central bank projections. Experimental Economics. Forthcoming.

34. Mitchell, J., & Weale, M. (2019). Forecasting with Unknown Unknowns: Censoring and Fat Tails on the Bank of England's Monetary Policy Committee (No. 27). Economic Modelling and Forecasting Group.

35. Neely, C. (2005). Using implied volatility to measure uncertainty about interest rates. Federal Reserve Bank of St. Louis Review, 87(3), 407-25.

36. Petersen, L. & Rholes, R. (2020). Escaping Secular Stagnation with Unconventional Policy.

37. Pfajfar, D., & Santoro, E. (2010). Heterogeneity, learning and information stickiness in inflation expectations. Journal of Economic Behavior & Organization, 75(3), 426-444.

38. Pfajfar, D., & Žakelj, B. (2014). Experimental evidence on inflation expectation formation. Journal of Economic Dynamics and Control, 44, 147-168.

39. Pfajfar, D., & Žakelj, B. (2016). Uncertainty in forecasting inflation and monetary policy design: Evidence from the laboratory. International Journal of Forecasting, 32(3), 849-864.

40. Pfajfar, D., & Žakelj, B. (2018). Inflation expectations and monetary policy design: Evidence from the laboratory. Macroeconomic Dynamics, 22(4), 1035-1075.

41. Rholes, R. & Sekhposyan, T. (2019). Central Bank Density Forecasts and Asset Prices: Do Revisions to Higher-order Moments Matter?

42. Rudebusch, G. D., & Williams, J. C. (2008). Revealing the secrets of the temple: The value of publishing central bank interest rate projections. In Asset Prices and Monetary Policy (pp. 247-289). University of Chicago Press.

43. Swanson, E. T. (2006). Have increases in Federal Reserve transparency improved private sector interest rate forecasts?. Journal of Money, Credit, and Banking, 38(3), 791-819.

44. Woodford, M. (2005). Central bank communication and policy effectiveness (No. w11898). National Bureau of Economic Research.

5. SUMMARY AND CONCLUSIONS

This dissertation presents three projects focused on central banking. In particular, this dissertation aims to understand emerging practices in central bank communication and to explore unconventional monetary policy alternatives. Together, these three chapters constitute my effort to better understand some of the primary tools available to policymakers who find themselves constrained at or near the effective lower bound and in need of alternative methods of stimulating economic activity.

In chapter one of this dissertation, I introduce a flexible experimental framework that serves as a testbed for monetary policy, fiscal policy, and central bank communication. Crucially, this framework provides a way for academics and policymakers alike to generate much-needed empirical data to guide real-world policy decisions. Though the laboratory necessarily abstracts away from some real-world frictions, this framework provides a low-cost, low-risk way for researchers to explore policy alternatives under tightly-controlled conditions. We then use this framework to investigate the ability of unconventional monetary policies to alleviate deflationary traps and achieve liftoff at the zero lower bound. We find that central bank intervention into deflationary traps via inflation targeting consistently fails to close output and employment gaps and to align aggregate inflation with the central bank's inflation target. This result is consistent with Arifovic and Petersen (2017), who find that neither qualitative nor quantitative communication of higher inflation targets in a liquidity trap is sufficient to stimulate inflation expectations in a learning-to-forecast experimental environment. Removing the zero lower bound and allowing interest rates to become negative, on the other hand, reliably aligns aggregate inflation in our experimental economies with the central bank's targeted full-employment, inflationary equilibria. Our results suggest that policies aimed at stimulating aggregate demand through increased real wealth balances are more effective than those relying on rational expectations.

Chapter two considers the BoE's density forecasts and its revisions to quantify the effects of information flow on the financial markets and expectations. The paper contributes to the broader literature on the effects of news on financial markets and to the emerging literature on how information in the higher-order moments of central bank forecasts matters from a policy perspective. We find that financial markets respond at least as strongly to the information contained in the higher-order moments of the BoE's density forecasts of output growth and inflation as to the information contained in the corresponding first moments. further, we find that both the magnitude and direction of responses are state-contingent. Additionally, we use Blue Chip Financial Forecast data to study how professional forecasters respond to information contained in higher-order forecast moments. We find that the consensus forecast and level of forecast disagreement of both short- and long-term interest rates are strongly correlated with higher-order forecast moments. Further, we observe a logical coherence between private expectations and realized yield changes, suggesting that market participants understand how rates will respond to higher-order forecast moments. Overall, our results suggest that communicating high-order forecasts moments to market participants does affect their subsequent behavior. Important from a policy perspective is that higher-order moments can move markets, which suggests that density forecasting is a viable policy option.

ur work aims to fill this gap by providing original evidence on the effects of communicating uncertainty on expectation formation. First, we study the introduction of a measure of a central bank's forecast uncertainty into central bank projections (i.e. the publication of density rather than point forecasts). Our interest is in how this affects aggregate dynamics and how forecasters incorporate information in the first and second central moments into their own expectations and perceptions of future uncertainty. Second, this paper studies behavior in individual-choice and coordination settings to understand the extent to which strategic concerns influence how agents use information when forming expectations.

We find that both point and density projections significantly improve forecast accuracy and decrease cross-sectional disagreement relative to an environment with no auxiliary central bank communication. This is consistent with empirical evidence that central bank projections coordinate expectations and reduce forecast errors. Furthermore, projections increase the proportion of participants who form ex-ante rational expectations. We provide new evidence showing that a large majority of participants use the same heuristics to formulate both their one- and two-period-ahead forecasts. However, roughly 20% of participants employ different heuristics when forecasting at different horizons. These subjects tend to use more irrational heuristics for their further ahead forecasts. However, projections nudge more distant forecasts toward the ex-ante rational prediction.

These results are in line with Mokhtarzadeh and Petersen (2020) who find inflation projections works effectively to guide expectations in the absence of a zero lower bound (ZLB). Inflation expectations are not as well managed by rationally–constructed inflation projections in the presence of ZLB. See Arifovic and Petersen (2017), Ahrens et al. (2018), and Kryvtsov and Petersen (2020) for examples of relatively poorer management of inflation expectations through projections at ZLB. In particular, more simplistic communications are more effective to guide expectations. Kryvtsov and Petersen note that simple, relatable information about past interest rates work more effectively to manage expectations than forward-looking projections about policy rates and forward guidance. Ahrens et al. find that gradual adjustment of inflation projections by human central bankers can more effectively build up credibility and manage expectations at the ZLB. Arifovic and Petersen find that qualitative rather than quantitative communication can work somewhat better to reduce pessimism.

Chapter three of this dissertation uses a 'learning to forecast' experiment based around a three-equation New Keynesian model to study the effect of introducing uncertaitny into central bank communications. We find that communicating uncertainty via density bands surrounding a point projection mutes the positive effects of publishing point projections. Compared to point projections, density projections significantly increase forecast errors and disagreement. The central bank transmits their uncertainty to forecasters, leading to higher levels of private forecast uncertainty. Moreover, fewer subjects form ex-ante rational expectations at both forecasting horizons. Credibility in the central bank's projections significantly decreases when participants are presented with a less precise projection. This result is in line with Baeriswyl and Cornand (2016) who show in a Keynesian beauty contest environment that subjects place more weight on public signals the more precise is the signal. Their results are more nuanced. Subjects overreact relative to theoretical predictions when the public signal is imprecise, and under-react when it is more precise. A notable finding in our experiment is that inflation volatility and the heterogeneity in forecasts and heuristics increases when participants have more market power in the Individual treatments. Other LTF experiments have also explored the effects of individual subjects' market size on system stability. Kopányi et al. (2019) show in an asset market LTF that increasing the market power of more accurate forecasters can lead to greater instability. Bao et al. (2020), on the other hand, find that asset price bubbles grow even faster with larger group sizes (where individuals have less market power) and participants are more likely to coordinate on trend-chasing strategies. The differences in our experimental findings and Bao et al. likely are driven by the relatively greater negative feedback present in our data-generating process, that encourage coordination on more stable heuristics (see Heemeijer et al. 2009 for evidence on the effects of positive and negative feedback in LTF experiments). We show that central bank projections both provide valuable information to reduce forecasters' confusion and can alleviate some strategic uncertainty. Our results line up with Akiyama et al. (2017) who show that mispricing commonly observed in experimental asset markets is driven at least in part by strategic uncertainty. In fact, they show that strategic uncertainty explains at least as much of median initial forecast deviation from the fundamental value as does confusion. Our results are consistent with the broader literature that shows individuals will anchor on simple and salient information when facing cognitive overload (Deck and Jahedi, 2015, Kryvtsov and Petersen, 2020). Much of the macroeconomic literature uses a notion of rational expectations that focuses on an agents' point forecast of some future event. There is scope in learning-to-forecast experiments to better understand participants' subjective uncertainty, how it relates to rationality, and how the relationship between rationality and uncertainty can be influenced by monetary policy and central bank communication. We find that participants exhibit an unusually high level of confidence in their own forecasts, and this confidence can be better strengthened through more precise information. An important avenue of future study is whether participants would act on expectations given their level of confidence.

APPENDIX A

A.1 Instructions

A.1.1 Group Instructions

EXPERIMENTAL STUDY OF ECONOMIC DECISION MAKING

Welcome! You are here to participate in an economic experiment. If you read these instructions carefully and make appropriate decisions, you may earn a considerable amount of money. We will pay you this money in cash immediately after this experiment.

Each of you will earn $10 for attending. This is your show-up fee. Throughout this experiment you will also earn points based on the decisions you make. Every point you earn is worth an additional $0.50. We reserve the right to improve the show up fee in your favour if average payoffs are lower than expected.

During the experiment you are not allowed to communicate with other participants. Please raise your hand if you have any questions. An experimenter will answer your questions privately. You will be excluded from the experiment and deprived of all payments aside from the show-up fee if you do not comply with these instructions.

This experiment is based on a simple simulation that approximates fluctuations in a real economy. Your task is to serve as private forecasters and provide real–time forecasts about future inflation in this simulated economy. These instructions will explain what inflation and the interest rate are, how they move around in this economy, and how they depend on your forecasts. We will allow you to practice making forecasts for several unpaid periods before we begin paid periods in this experiment. You will then participate in two sequences of 30 paid periods, for a total of 60 paid periods of play.

In this simulation, households and firms (whose decisions are automated by the computer) will form forecasts identically to yours. So to some degree, outcomes that you will see in the game will depend on the way in which all of you form your forecasts. However, your earnings in this experiment depend on the accuracy of your individual forecasts.

You will also submit a measure of uncertainty about your forecast called your anticipated forecasting error. You will earn money if actual inflation is within the bounds of this error. Otherwise, you will earn nothing.

Please note that all values are given in basis points, a measurement often used in descriptions of the economy. All values can be positive, negative, or zero at any point in time.

Overview of the Economy

In each period, you will submit a forecast of inflation for the next two periods. For example, suppose it is now period 10. Then you will submit a forecast of inflation in period 11 and a forecast of inflation in period 12. By 'forecast of inflation' we mean your best guess of what inflation will be. The more accurate your guess, the more money you will earn.

Your forecasts should be given in basis points. Here are some examples of the relationship between basis points and percentages:

1% = 100 basis points

3.25% = 325 basis points

-0.5% = -50 basis points

-4.8% = -480 basis points

You can submit any forecast you wish, positive or negative or zero, but please only submit integers.

The economy consists of three main variables:

- **Inflation** – Inflation is the change in price that occurs between two periods.
- **Interest Rates** – The interest rate is the amount of money that people earn on savings. A higher interest rate entices consumers to save more and spend less on consumption. Thus, a higher interest rate puts downward pressure on inflation.
- **Shocks** – Shocks are changes to the amount consumers in the economy wish to purchase. Shocks change every period and are influenced by a random component and by past shocks. A positive shock today increases inflation today and vice versa.

Your goal in this experiment is to forecast future inflation as accurately as possible. Thus, we now provide detailed explanations of the factors that influence inflation and the relationships between the different variables in the economy.

Shocks:

Intuitively, you can think of shocks as weather shocks. Over the long run, the weather has no effect on how much consumers want to buy. However, from day to day, there may be random changes to the weather that do influence what people do and buy. You can think of a positive shock as unexpectedly nice weather. When the weather is especially nice, consumers are spending more time out of their homes and increasing their expenditures (for example, buying ice cream, going out for a nice dinner, or going to the beach). A negative shock can be thought of as unexpectedly terrible weather. This bad weather makes it so that people do now want to leave their homes, causing expenditures to be relatively low. Gradually, the shocks, like weather, will revert back to their long-run levels. As the shocks dissipate, new random events occur that will make consumers want to increase or decrease their spending. Shocks will have a precise value and will be displayed on your screen.

Whenever a positive shock occurs and spending increases, this will put upward pressure on prices (i.e. upward pressure on inflation). Conversely, a negative shock will put downward pressure on prices (i.e. downward pressure on inflation).

We calculate the values of a shock in each period as follows:

$$Shock_p = 0.57(Shock_{p-1}) + RandomComponent_p$$

- The random component is 0 on average
- Roughly two out of three times the shock will be between -138 and 138 basis points.

- 95% of the time the shock will be between -276 and 276 basis points

For example, shocks may evolve as follows:

$$
\begin{aligned}
Shock_1 &= 30 \\
Shock_2 &= 30 \times 0.57 + New\ Draw \\
&= 17.1 + (newdraw) \\
Shock_2 &= 17.1 + (-150) \\
&= -132.9 \\
Shock_3 &= -132.9 \times .57 + New\ Draw \\
&= ...
\end{aligned}
$$

Interest Rates: The central bank in this economy will adjust the nominal interest rate in each period to keep inflation as close to zero as possible. As inflation increases, the central bank will increase the nominal interest rate more than one-for-one with inflation. An increase in the nominal interest rate has a direct negative effect on consumer demand and production, and an indirect negative effect on inflation. Importantly, you will not observe the current interest rate when you are forming your inflation forecasts. After you submit your forecasts, the computer will solve for the current period's inflation using the median forecasts from all subjects in the room and the current-period shock shock (which you will see). It is important for you to realize that, even though the central bank is aiming for zero inflation, it will rarely accomplish this. This is because of the random shocks that occur in each period and the public's expectations. However, the central bank will keep the economy more stable than the economy would be in the absence of the central bank.

How the economy evolves:

Each period, you and the other forecasters in this room will submit your beliefs about inflation for the next period and the period after that. To be clear, if we are in period 10, you will submit an inflation forecast for period 11 and for period 12. The software will select the median of each of the two forecasts as the aggregate forecasts. The software uses the median, rather than the average forecast, so that a small number of subjects cannot have a significant effect on the economy.

These aggregate forecasts play an important role in determining inflation today. This is because inflation today is determined largely by aggregate forecasts about future inflation. If the majority of forecasters expect relatively high inflation tomorrow, then inflation today will be higher. The idea behind this is simple: If the professional forecasters communicate to the public that inflation is likely to rise tomorrow, consumers will spend more immediately to avoid paying the relatively higher prices tomorrow. This increase in demand today will cause prices to start rising today, and so inflation will increase today. Likewise, if the median forecaster predicts higher inflation for two days from now, households will need to have a bit more money tomorrow than they would otherwise to avoid paying the higher prices predicted for two days from now.

More precisely, inflation and interest rates evolve according to the following equations:

$$\begin{aligned} Inflation_t &= 1.54(\text{Median forecast of Inflation}_{t+1}) - 0.58(\text{Median forecast of Inflation}_{t+2}) \\ &\quad +0.08(\text{Shock}_t) \end{aligned}$$

$$\text{Interest Rate}_t = 4.44(\text{Median forecast of Inflation}_{t+1}) - 3.12(\text{Median forecast of Inflation}_{t+2})$$
$$+\ 0.41(\text{Shock}_t)$$

Important information about this economy:

- The Central Bank sets the target inflation at zero. In order to achieve this target it will adjust the nominal interest rate in each period. In some cases the nominal interest rate can become negative.

- Expectations about tomorrow (if in period 10, this is your forecast for period 11) are self-fulfilling in this economy. If you forecast higher inflation tomorrow then inflation will grow higher in the current period. Similarly, a median forecast of lower inflation tomorrow will cause inflation to fall in the current period.

- Expectations about two days from now (if in period 10, this is your forecast for period 12) relate negatively to inflation today. If you forecast higher inflation for two-days from now, then inflation today will fall. If instead you forecast lower inflation for two days from now, inflation today will increase.

Score

Your forecasting score in each period will depend on the accuracy of the forecasts you formed in the previous two periods. At the end of each period, the software will evaluate how accurate your forecasts from one- and two-periods ago were about the inflation rate in the current period. The difference between these numbers forms your absolute forecast error. The larger this absolute error, the lower is your forecasting score in that period. The letter p in the following example stands for 'period'.

- Absolute Forecast Error = $\|$Your Forecast $-$ Actual Value$\|$

- Total Score$_p$ =0.3($2^{-AbsoluteForecastError_{p-1}} + 2^{-AbsoluteForecastError_{p-2}}$)

The maximum score you can earn for forecasting in each period is 0.60 points. Your score will decrease exponentially as your forecast error increases. Suppose your forecast errors for inflation is:

0: Your score will be 0.6

50: Your score will be 0.42

100: Your score will be 0.30

200: Your score will be 0.15

300: Your score will be 0.075

500: Your score will be 0.02

1000: Your score will be 0

2000: Your score will be 0

Making decisions in this experiment

During this experiment, your main screen will display information that will help you make forecasts and earn more points.

At the top left of the screen, you will see your subject number, the current period, time remaining, and the total number of points you've earned through the previous period. You will also see three history plots on your screen.

The top history plot displays past interest rates and current and past shocks.

The second history plot shows your 1-period-ahead points forecasts of inflation (blue dots), error bands that you create with your anticipated forecasting error (blue shading centered around your point forecasts) and actual inflation (red dots). Note that the difference between your forecasts of one-period-ahead inflation (blue dots) and the actual levels of inflation (red dots) constitutes your one-period-ahead forecast error in past periods.

The third history plot shows your 2-period-ahead point forecasts of inflation (orange dots), error bands that you create with your anticipated forecasting error (orange shading centered around your point forecasts), and actual inflation (red dots). Note that the difference between your forecasts of two-period-ahead inflation (orange dots) and the actual levels of inflation (red dots) constitute your two-period-ahead forecast error in past periods.

Note: this section read one of three ways depending upon treatment:

For NoComm, skip directly to "You have 65 seconds...

For Point treatments:

Both the second and third plots also contain the Central Bank's forecast of inflation for the next five periods (green). It is important to remember that the projections are simply a forecast and not a promise. The Central Bank uses the model discussed earlier in these instructions, and the current and expected future shocks, to form its projections. In particular, it predicts that the economy will return to zero levels of inflation in the near future.

For Point&Density treatments:
Both the second and third plots also contain the Central Bank's forecast of inflation for the next five periods (green). This forecasts also includes green shading, which represents the Central Bank's level of uncertainty for its corresponding point projections. These bands will contain the correct realization of inflation about 66% of the time. It is important to remember that the projections are simply a forecast and not a promise. The Central Bank uses the model discussed earlier in these instructions, and the current and expected future shocks, to form its projections. In particular, it predicts that the economy will return to zero levels of inflation in the near future.

You have 65 seconds to make decisions in the first nine periods and only 50 seconds thereafter. You may submit both negative and positive forecasts and forecasts of 0. Please review your forecasts before pressing the SUBMIT button because you cannot revise your forecasts afterward.

The anticipated forecast error:
You must also submit a measure of how uncertain you are about your inflation forecasts. We call this your anticipated forecasting error.Note this value should always be positive and your error bounds are centered around your point forecast.

Suppose you forecast inflation tomorrow to be 10 basis points but feel more confident that actual inflation will fall between 5 and 15 basis points. You should indicate this by submitting an anticipated forecasting error of 5. This forms anticipated error bounds of 5 to 15 since $10 - 5 = 5$ and $10 + 5 = 15$. If actual inflation is any number from 5 to 15, we pay you. Otherwise, you earn nothing for this anticipated forecasting error.

If actual inflation falls within your anticipated forecast error bounds, then we pay your anticipated forecast error according to the following function:

$$AnticipatedErrorEarnings = \frac{15}{10+anticipated\ error}$$

Notice that your earnings for your anticipated forecast error decrease as your anticipated forecast error increases. However, it is important for you to understand that we pay you this amount ONLY if the realized value of inflation lies inside your anticipated forecasting error bands. If actual inflation is outside your anticipated forecasting error bands, then you earn 0 points for providing your anticipated forecasting error.

An example: Suppose it is period 3. Suppose in periods 1 and 2 you provided an inflation forecast of 10 basis points for period 3 inflation. Suppose your anticipated forecasting error in period 1 was 5 and in period 2 it was 10. Then your error bounds for period 1 are 5 to 15 and for period 2 are 0 to 20. Suppose actual inflation at the end of period 3 is 17. Then you earn 0 points for your anticipated forecast error provided in period 1. This is because 17 is not between 5 and 15. However, you would earn $\frac{15}{10+10} = .75$ points for your anticipated forecast error provided in period 2, since 17 is between 0 and 20.

Our software will randomly select (with equal probability) to pay you for **either** your point forecasts **or** for your anticipated forecast error in each period of play. **We will never pay for both in a single period**.

A.1.2 Individual Instructions

EXPERIMENTAL STUDY OF ECONOMIC DECISION MAKING

Welcome! You are here to participate in an economic experiment. If you read these instructions carefully and make appropriate decisions, you may earn a considerable amount of money. We will pay you this money in cash immediately after this experiment.

Each of you will earn $10 for attending. This is your show-up fee. Throughout this experiment you will also earn points based on the decisions you make. Every point you earn is worth an additional $0.50. We reserve the right to improve the show up fee in your favour if average payoffs are lower than expected.

During the experiment you are not allowed to communicate with other participants. Please raise your hand if you have any questions. An experimenter will answer your questions privately. You will be excluded from the experiment and deprived of all payments aside from the show-up fee if you do not comply with these instructions.

This experiment is based on a simple simulation that approximates fluctuations in a real economy. Your task is to serve as the sole private forecaster in your own economy and provide real–time forecasts about future inflation in this simulated economy. These instructions will explain what inflation and the interest rate are, how they move around in this economy, and how they depend on your forecasts. We will allow you to practice making forecasts for several unpaid periods before we begin paid periods in this experiment. You will then participate in two sequences of 30 paid periods, for a total of 60 paid periods of play.

In this simulation, households and firms (whose decisions are automated by the computer) will form forecasts identically to yours. Your earnings in this experiment depend on the accuracy of your individual forecasts.

You will also submit a measure of uncertainty about your forecast called your anticipated forecasting error. You will earn money if actual inflation is within the bounds of this error. Otherwise, you will earn nothing.

Please note that all values are given in basis points, a measurement often used in descriptions of the economy. All values can be positive, negative, or zero at any point in time.

Overview of the Economy

In each period, you will submit a forecast of inflation for the next two periods. For example, suppose it is now period 10. Then you will submit a forecast of inflation in period 11 and a forecast of inflation in period 12. By 'forecast of inflation' we mean your best guess of what inflation will be. The more accurate your guess, the more money you will earn.

Your forecasts should be given in basis points. Here are some examples of the relationship between basis points and percentages:

1% = 100 basis points

3.25% = 325 basis points

-0.5% = -50 basis points

-4.8% = -480 basis points

You can submit any forecast you wish, positive or negative or zero, but please only submit integers.

The economy consists of three main variables:

- **Inflation** – Inflation is the change in price that occurs between two periods.
- **Interest Rates** – The interest rate is the amount of money that people earn on savings. A higher interest rate entices consumers to save more and spend less on consumption. Thus, a higher interest rate puts downward pressure on inflation.
- **Shocks** – Shocks are changes to the amount consumers in the economy wish to purchase. Shocks change every period and are influenced by a random component and by past shocks. A positive shock today increases inflation today and vice versa.

Your goal in this experiment is to forecast future inflation as accurately as possible. Thus, we now provide detailed explanations of the factors that influence inflation and the relationships between the different variables in the economy.

Shocks:
Intuitively, you can think of shocks as weather shocks. Over the long run, the weather has no effect on how much consumers want to buy. However, from day to day, there may be random changes to the weather that do influence what people do and buy. You can think of a positive shock as unexpectedly nice weather. When the weather is especially nice, consumers are spending more time out of their homes and increasing their expenditures (for example, buying ice cream, going out for a nice dinner, or going to the beach). A negative shock can be thought of as unexpectedly terrible weather. This bad weather makes it so that people do now want to leave their homes, causing expenditures to be relatively low. Gradually, the shocks, like weather, will revert back to their long-run levels. As the shocks dissipate, new random events occur that will make consumers want to increase or decrease their spending. Shocks will have a precise value and will be displayed on

your screen.

Whenever a positive shock occurs and spending increases, this will put upward pressure on prices (i.e. upward pressure on inflation). Conversely, a negative shock will put downward pressure on prices (i.e. downward pressure on inflation).

We calculate the values of a shock in each period as follows:

$$Shock_p = 0.57(Shock_{p-1}) + RandomComponent_p$$

- The random component is 0 on average
- Roughly two out of three times the shock will be between -138 and 138 basis points.
- 95% of the time the shock will be between -276 and 276 basis points

For example, shocks may evolve as follows:

$$\begin{aligned}
Shock_1 &= 30 \\
Shock_2 &= 30 \times 0.57 + New\ Draw \\
&= 17.1 + (newdraw) \\
Shock_2 &= 17.1 + (-150) \\
&= -132.9 \\
Shock_3 &= -132.9 \times .57 + New\ Draw \\
&= ...
\end{aligned}$$

Interest Rates: The central bank in this economy will adjust the nominal interest rate in each period to keep inflation as close to zero as possible. As inflation increases, the central bank will increase the nominal interest rate more than one-for-one with inflation. An increase in the nominal interest rate has a direct negative effect on consumer demand and production, and an indirect negative effect on inflation. Importantly, you will not observe the current interest rate when you are forming your inflation forecasts. After you submit your forecasts, the computer will solve for the current period's inflation using your forecasts and the current-period shock (which you will see). It is important for you to realize that, even though the central bank is aiming for zero inflation, it will rarely accomplish this. This is because of the random shocks that occur in each period and the public's expectations. However, the central bank will keep the economy more stable than the economy would be in the absence of the central bank.

How the economy evolves:

Each period, you will submit your beliefs about inflation for the next period and the period after that. To be clear, if we are in period 10, you will submit an inflation forecast for period 11 and for period 12. The software will use your two forecasts as the aggregate forecasts.

These aggregate forecasts play an important role in determining inflation today. This is because inflation today is determined largely by aggregate forecasts about future inflation. If you expect relatively high inflation tomorrow, then inflation today will be higher. The idea behind this is simple: If you communicates to the public that inflation is likely to rise tomorrow, consumers will spend more immediately to avoid paying the relatively higher prices tomorrow. This increase in demand today will cause prices to start rising today, and so inflation will increase today. Likewise, if you predict higher inflation for two days from now, households will need to have a bit more money tomorrow than they would otherwise to avoid paying the higher prices predicted for two days from now.

More precisely, inflation and interest rates evolve according to the following equations:

$$\begin{aligned} Inflation_t &= 1.54(\text{Your forecast of Inflation}_{t+1}) - 0.58(\text{Median forecast of Inflation}_{t+2}) \\ &\quad +0.08(\text{Shock}_t) \end{aligned}$$

$$\text{Interest Rate}_t = 4.44(\text{Your forecast of Inflation}_{t+1}) - 3.12(\text{Median forecast of Inflation}_{t+2}) + 0.41(\text{Shock}_t)$$

Important information about this economy:

- The Central Bank sets the target inflation at zero. In order to achieve this target it will adjust the nominal interest rate in each period. In some cases the nominal interest rate can become negative.
- Expectations about tomorrow (if in period 10, this is your forecast for period 11) are self-fulfilling in this economy. If you forecast higher inflation tomorrow then inflation will grow higher in the current period. Similarly, a median forecast of lower inflation tomorrow will cause inflation to fall in the current period.
- Expectations about two days from now (if in period 10, this is your forecast for period 12) relate negatively to inflation today. If you forecast higher inflation for two-days from now, then inflation today will fall. If instead you forecast lower inflation for two days from now, inflation today will increase.

Score

Your forecasting score in each period will depend on the accuracy of forecasts formed in the previous two periods. At the end of each period, the software will evaluate how accurate your forecasts from one- and two-periods ago were about the inflation rate in the current period. The difference between these numbers forms your absolute forecast error. The larger this absolute error, the lower is your forecasting score in that period. The letter p in the following example stands for 'period'.

- Absolute Forecast Error = $\|$Your Forecast $-$ Actual Value$\|$

- Total Score$_p$ =0.3$(2^{-AbsoluteForecastError_{p-1}} + 2^{-AbsoluteForecastError_{p-2}})$

The maximum score you can earn for forecasting in each period is 0.60 points. Your score will decrease exponentially as your forecast error increases. Suppose your forecast errors for inflation is:

0: Your score will be 0.6

50: Your score will be 0.42

100: Your score will be 0.30

200: Your score will be 0.15

300: Your score will be 0.075

500: Your score will be 0.02

1000: Your score will be 0

2000: Your score will be 0

Making decisions in this experiment

During this experiment, your main screen will display information that will help you make forecasts and earn more points.

At the top left of the screen, you will see your subject number, the current period, time remaining, and the total number of points you've earned through the previous period. You will also see three history plots on your screen.

The top history plot displays past interest rates and current and past shocks.

The second history plot shows your 1-period-ahead points forecasts of inflation (blue dots), error bands that you create with your anticipated forecasting error (blue shading centered around your point forecasts) and actual inflation (red dots). Note that the difference between your forecasts of one-period-ahead inflation (blue dots) and the actual levels of inflation (red dots) constitutes your one-period-ahead forecast error in past periods.

The third history plot shows your 2-period-ahead point forecasts of inflation (orange dots), error bands that you create with your anticipated forecasting error (orange shading centered around your point forecasts), and actual inflation (red dots). Note that the difference between your forecasts of two-period-ahead inflation (orange dots) and the actual levels of inflation (red dots) constitute your two-period-ahead forecast error in past periods.

Note: this section read one of three ways depending upon treatment:

For NoComm, skip directly to "You have 65 seconds...

For Point treatments:
Both the second and third plots also contain the Central Bank's forecast of inflation for the next five periods (green). It is important to remember that the projections are simply a forecast and not a promise. The Central Bank uses the model discussed earlier in these instructions, and the current and expected future shocks, to form its projections. In particular, it predicts that the economy will return to zero levels of inflation in the near future.

For Point&Density treatments:
Both the second and third plots also contain the Central Bank's forecast of inflation for the next five periods (green). These forecasts also includes green shading, which represents the Central Bank's level of uncertainty for its corresponding point projections. These bands will contain the correct realization of inflation about 66% of the time. It is important to remember that the projections are simply a forecast and not a promise. The Central Bank uses the model discussed earlier in these instructions, and the current and expected future shocks, to form its projections. In particular, it predicts that the economy will return to zero levels of inflation in the near future.

You have 65 seconds to make decisions in the first nine periods and only 50 seconds thereafter. You may submit both negative and positive forecasts and forecasts of 0. Please review your forecasts before pressing the SUBMIT button because you cannot revise your forecasts afterward.

The anticipated forecast error:
You must also submit a measure of how uncertain you are about your inflation forecasts. We call this your anticipated forecasting error.Note this value should always be positive and your error bounds are centered around your point forecast.

Suppose you forecast inflation tomorrow to be 10 basis points but feel more confident that actual inflation will fall between 5 and 15 basis points. You should indicate this by submitting an anticipated forecasting error of 5. This forms anticipated error bounds of 5 to 15 since $10 - 5 = 5$ and $10 + 5 = 15$. If actual inflation is any number from 5 to 15, we pay you. Otherwise, you earn nothing for this anticipated forecasting error.

If actual inflation falls within your anticipated forecast error bounds, then we pay your anticipated forecast error according to the following function:

$$AnticipatedErrorEarnings = \frac{15}{10 + anticipated\ error}$$

Notice that your earnings for your anticipated forecast error decrease as your anticipated forecast error increases. However, it is important for you to understand that we pay you this amount ONLY if the realized value of inflation lies inside your anticipated forecasting error bands. If actual inflation is outside your anticipated forecasting error bands, then you earn 0 points for providing your anticipated forecasting error.

An example: Suppose it is period 3. Suppose in periods 1 and 2 you provided an inflation forecast of 10 basis points for period 3 inflation. Suppose your anticipated forecasting error in period 1 was 5 and in period 2 it was 10. Then your error bounds for period 1 are 5 to 15 and for period 2 are 0 to 20. Suppose actual inflation at the end of period 3 is 17. Then you earn 0 points for your anticipated forecast error provided in period 1. This is because 17 is not between 5 and 15. However, you would earn $\frac{15}{10+10} = .75$ points for your anticipated forecast error provided in period 2, since 17 is between 0 and 20.

Our software will randomly select (with equal probability) to pay you for **either** your point forecasts **or** for your anticipated forecast error in each period of play. **We will never pay for both in a single period**.

Figure A.1: Aggregate statistics in two- and one-period ahead NK DGPs

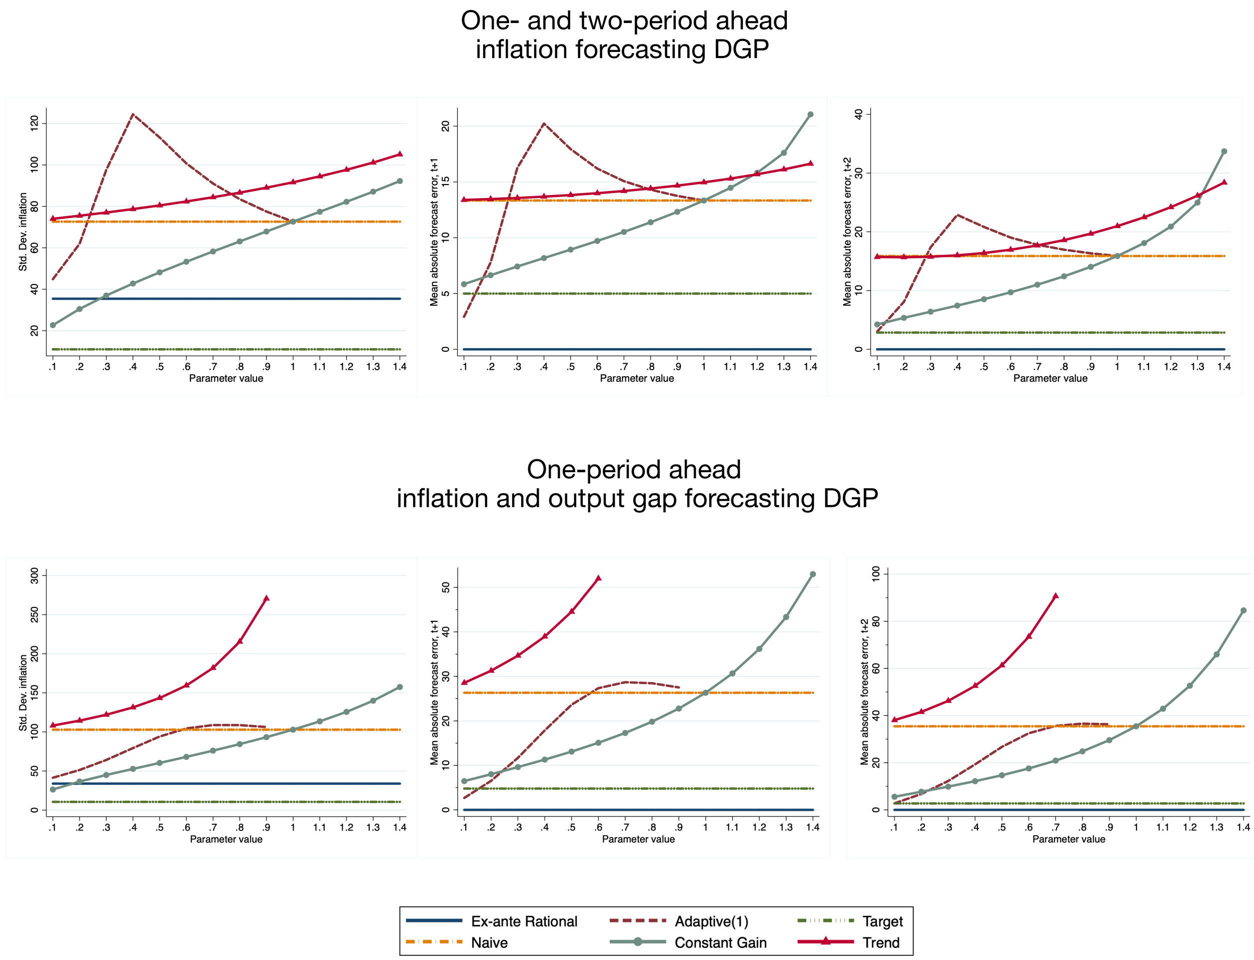

A.2 Comparison of two-period ahead and one-period ahead data-generating processes

The linearized New Keynesian model is typically presented as a function of one-period ahead inflation and output gap expectations. The model can be rederived as a function of one- and two-period ahead inflation expectations. We show in Figure Figure A.1 the impact of this modeling choice on the standard deviation of inflation and one- and two-period ahead absolute inflation forecast errors under alternative forecasting heuristics under alternative parameterizations of the Adaptive(1), Constant Gain, and Trend-Chasing models (α, γ, and τ). The top row presents simulated results for the version that we implement in our experiment (where participants submit one- and two-period ahead inflation forecasts), while the bottom row presents the results for the more standard version (where participants submit one-period ahead inflation and output gap forecasts). For simplicity, we will describe these as two- and one-period ahead models, respectively. To economize on space, we only include inflation expectations. Note that the one-period ahead model also employs output gap expectations. For comparison, we also calculate and include the two-period ahead forecast errors in the more conventional framework. The models are simulated for 50,000 periods and approximated moments are calculated.

Under rational expectations, these models generate nearly identical simulations. Mean absolute forecast errors are zero in both cases. The standard deviations of inflation in the two-period ahead model are 35.41 and 33.90 bps in the two- and one-period ahead models, respectively. They differ only because of the rounding that we employed in our experimental design (up to two decimal points).

Under Adaptive(1) expectations, we assume that agents form the following inflation expectations:

$$E_t\pi_{t+1} = \alpha\pi_{t-1} + (1-\alpha)\pi_{t+1}$$
$$E_t\pi_{t+2} = \alpha\pi_{t-1} + (1-\alpha)\pi_{t+2}$$

and output gap expectations in a similar manner. In both the two- and one-period ahead models, we see that the standard deviations of inflation are rather comparable, ranging from 45 bps to 113 bps in the two-period ahead model and 41 bps to 108 bps in the one-period ahead model. Forecast errors are also quite comparable. Absolute one-period ahead inflation forecast errors range from 3 to 20 bps in the two-period ahead model and 3 to 28 bps in the one-period ahead model. Absolute two-period ahead inflation forecast errors range from 3 to 22 bps in the two-period ahead model and 3 to 27 bps in the one-period ahead model.

Under Target expectations, we assume that agents form the following inflation expectations:

$$E_t\pi_{t+1} = E_t\pi_{t+2} = 0$$

and the same for output gap expectations. The results are highly comparable across the two types of DGPs. Standard deviations are 11 and 10.5 bps in the two- and one-period ahead models. Likewise, absolute forecast errors are approximately the same (5 bps for one-period ahead forecasts and 3 bps for two-period ahead forecasts).

Under Naive expectations, we assume that agents simply use last period's outcomes to form their inflation expectations:

$$E_t\pi_{t+1} = E_t\pi_{t+2} = \pi_{t-1}.$$

Under Constant Gain expectations, agents are assumed to update their past forecasts in response to their most recently observed past forecast errors.

$$E_t\pi_{t+1} = E_{t-2}\pi_{t-1} - \gamma(E_{t-2}\pi_{t-1} - \pi_{t-1})$$
$$E_t\pi_{t+2} = E_{t-3}\pi_{t-1} - \gamma(E_{t-3}\pi_{t-1} - \pi_{t-1})$$

Finally, under Trend-Chasing expectations, agents are assumed to extrapolate recent trends when forming their forecasts:

$$E_t\pi_{t+1} = \pi_{t-1} + \tau(\pi_{t-1} - \pi_{t-2})$$
$$E_t\pi_{t+2} = \pi_{t-1} + 2\tau(\pi_{t-1} - \pi_{t-2})$$

Naive, Constant Gain, and especially Trend-Chasing heuristics lead to much less volatile results in our two-period ahead version of the New Keynesian model. For the Naive model, the standard deviation increases from 73 bps to 103 bps by moving from a two- to one-period ahead model. Increasing the constant gain parameter, γ, or the trend-chasing parameter, τ, leads to significantly more volatility in the one-period ahead model than in the two-period ahead model. Forecast errors follow a similar pattern.

The relative volatilities and forecast errors are mostly preserved as we change the DGP to a two-period ahead framework. One exception is with the Trend-Chasing heuristic, which generates relatively more stable dynamics and smaller forecast errors compared to Constant Gain as the parameters grow large. This is likely because of the negative feedback associated with $t + 2$ forecasts generated in our revised DGP. Extrapolative two-period ahead expectations would serve to mute inflation in our two-period ahead model but not in the one-period ahead model.

A.3 Time series

Below we present comparisons of Group and Individual inflation rates, by treatment, in Figure A.2 to Figure A.4. The thick dark line denotes inflation in the Group treatment, while the light lines indicate inflation for each individual subject's session in the Individual treatment.

Figure A.2: Time series of NoComm sessions

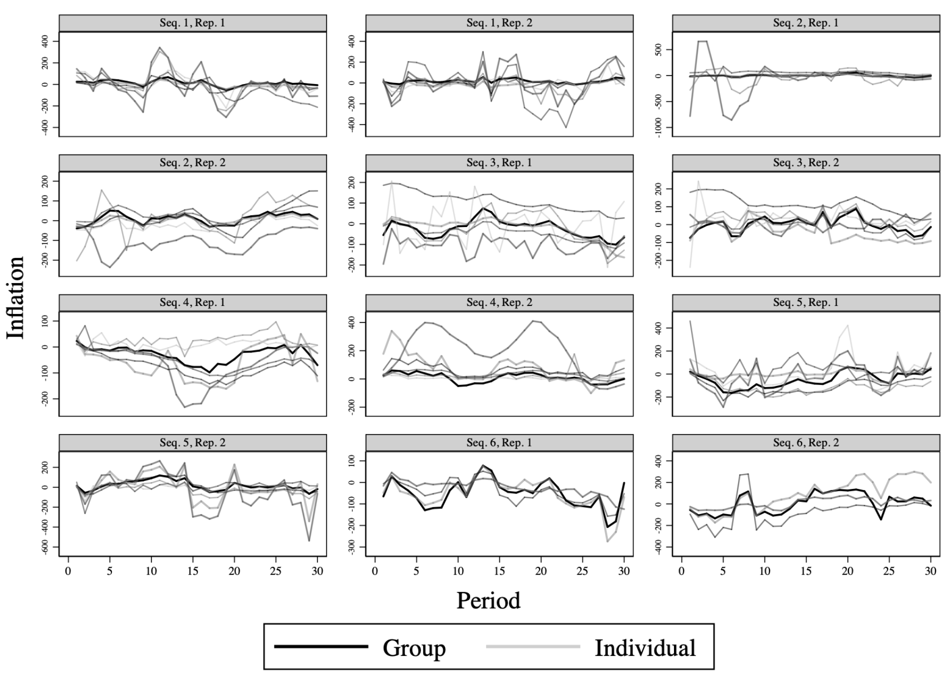

Figure A.3: Time series of Point sessions

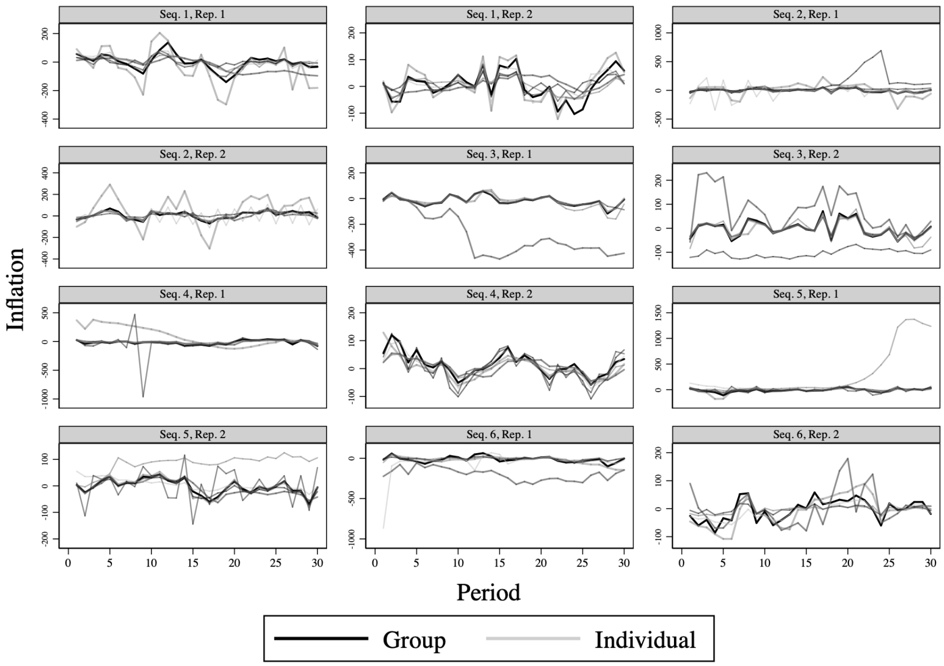

Figure A.4: Time series of Point&Density sessions