Price-Level Targeting and Inflation Expectations: Experimental Evidence*

Malik Shukayev[a] and Robert Amano[b]
[a]University of Alberta
[b]Pure Macro Research Inc. and Eco.org

The theoretical stabilization benefits of price-level targeting (PLT) in a low interest rate environment generated substantial interest in PLT among academics and policymakers. This paper uses a controlled laboratory environment to simulate a sudden transition from inflation targeting to PLT to assess whether people would change their inflation expectations to be consistent with the new policy framework. Without such expectation adjustments, PLT often leads to worse outcomes than inflation targeting in mainstream macroeconomic models. Our results suggest that participants formulate inflation expectations consistent with the target-reverting nature of the price level but do not fully utilize it in forecasting inflation.

JEL Codes: E32, E52.

1. Introduction

Owing to concerns that a lower neutral interest rate will leave the interest rate policy constrained by the zero lower bound more frequently, the Federal Reserve System announced a set of modifications in its monetary policy framework. One of the key new elements stated in the 2020 FOMC "Statement on Longer-Run Goals and Monetary Policy Strategy" is adherence to a flexible average inflation targeting (AIT) framework, which "seeks to achieve inflation

*Klaus Adam, Jean Boivin, Jim Engle-Warnick, Bob King, Oleksiy Kryvtsov, John Leahy, Gregor Smith, and Dave Wolf offered valuable comments and suggestions. We thank Jim Engle-Warnick and Julie Héroux for organizing and conducting the experiments reported in this paper. We also thank Jill Ainsworth, Pujan Thakrar, and Fang Li for excellent research assistance. Author e-mails: shukayev@ualberta.ca and robert@puremacro.com.

that averages 2 percent over time and therefore judges that, following periods when inflation has been running persistently below 2 percent, appropriate monetary policy will likely aim to achieve inflation moderately above 2 percent for some time.” Notably, the length of period over which the 2 percent average is defined and calculated remained unspecified. This lack of clarity regarding key implementation detail is likely to make it harder for people to understand the implications of this change for inflation dynamics. For example, a traditional inflation-targeting (IT) framework, practiced in many countries, aims to stabilize a year-over-year inflation rate, which can be thought of as an AIT with an annual averaging window. Coibion et al. (2023) use U.S. household survey evidence to show that most people did not understand the difference between IT and AIT. Moreover, even people who understood AIT had practically identical expectations regarding future inflation and other macroeconomic variables as households without such understanding. Based on these findings, Coibion et al. (2023) conclude, “Although economic theory predicts that AIT can typically generate better economic outcomes than IT when policy is constrained at the zero lower bound by committing to higher future inflation in order to make up for current or past downside misses, we find no evidence that real-world consumers see this mechanism at work.”

AIT and other similar *makeup* strategies, defined by Hebden et al. (2020) as “policies that aim to offset, at least in part, past misses of inflation from its objective,” are examples of *history-dependent* monetary policy frameworks, which condition the current short-term inflation objective on the history of past deviations of inflation away from the long-term inflation target. Price-level targeting (PLT) is a limiting case of makeup strategies in which the inflation rate is averaged over an indefinitely long period to keep the price level (index) stabilized around a predetermined time path, set and announced during initiation of the PLT framework. An important advantage of PLT over IT and AIT is that it limits the long-run uncertainty regarding the future range of possible price-level realizations, thus simplifying planning and investment into long-term nominal assets (Meh, Ríos-Rull, and Terajima 2010). However, most of the existing PLT studies, both from central banks and academia, focus on the advantages of a credible PLT over IT in terms of its short-term macroeconomic stabilization properties.

Under PLT, periods of lower-than-average inflation must be followed by periods of higher-than-average inflation in order to return the price level to its target path. When inflation falls during a recession, expectations of future inflation tend to be higher, which, in turn, provides additional stimulus to the economy in the form of lower real interest rates. The built-in stimulus implied by a credible PLT is likely to be particularly valuable when the nominal interest rates are near the zero lower bound (ZLB) and cannot be lowered much further. This perceived stabilization advantage of PLT over IT, when the policy rate is at or near zero, has been the main reason for a recent spike of attention to PLT among U.S. monetary policy officials (Bernanke 2017; Bullard 2018; Williams 2018).

Much academic work has explored the merits of price-level targeting at the zero bound. Eggertsson and Woodford (2003), Svensson (2003), and Evans (2010), among others, proposed PLT as a means of mitigating policy constraints near the ZLB. At the same time, Svensson (1999), Woodford (2003), and Vestin (2006) find PLT to perform better than IT in business cycle stabilization even when ZLB is not an issue. More recently, Coibion and Gorodnichenko (2011) show that in a canonical New Keynesian model with positive trend inflation, interest rate rules that respond to price-level fluctuations are more likely to satisfy the Taylor principle for dynamic stability than policy rules that respond to inflation fluctuations, and thus are more robust to equilibrium indeterminacy.

Importantly, however, the efficacy of PLT in stabilizing the economy hinges on a key assumption: inflation expectations must be generated rationally (in a Muth sense) consistent with the PLT regime. If agents do not form their expectations in such a way, then it is entirely possible for PLT to deliver results that are inferior to IT. Honkapohja and Mitra (2020) study global stability properties of IT and PLT policy rules under learning in a New Keynesian model with ZLB constraints and conclude: "The performance of PLT is clearly better than inflation targeting, provided private agents' learning has at least partly incorporated the guidance from the price level target path. If private agents' learning does not use the guidance at all, IT has a larger domain of attraction [of a desirable steady state] than PLT. Credibility of the PLT regime is important for its success." Given the prominence of PLT as an alternative monetary policy framework, an important policy question is whether private agents would, in fact, adjust inflation expectations in a manner consistent with the new policy regime following a transition from IT to PLT. Besides possible credibility issues (Cateau and Shukayev 2022), one immediate concern is whether people would actually understand the implications of the new policy framework for short-run inflation dynamics. The findings of Coibion et al. (2023) about limited AIT understanding suggest similar comprehension problems for PLT. However, it is likely that the PLT objective is easier to communicate because a fixed price-level target path is much simpler to visualize than a rolling AIT averaging window.

While there is no direct macroeconomic evidence regarding inflation expectations under PLT regimes,[1] there is an active line of experimental papers comparing PLT with other monetary policy regimes in a controlled laboratory environment (Arifovic and Petersen 2017; Hommes and Makarewicz 2021; Kostyshyna, Petersen, and Yang 2022; and Salle 2023). A common finding of these papers is that PLT does not offer many advantages in terms of its stabilization properties relative to IT or AIT with short averaging horizons. Participants of these studies often fail to coordinate their individual forecasts on a stable PLT equilibrium with a detrimental impact on economic stability.

Our paper uses a lab experiment to simulate a sudden transition from IT to PLT to test whether individuals in our study understand the implications of a monetary policy framework switch and change their inflation expectations accordingly. In contrast to previously mentioned experimental studies, we abstract from any stability of equilibrium issues by focusing on individual inflation forecasts, which have no impact on the simulated model. In our study, individual participants see outcomes from a simulated rational expectations equilibrium and are incentivized to forecast the inflation rate realizations in the model. There are advantages and disadvantages to our approach. The main advantage is that in the absence of any endogenous effects of individual forecasts on the simulated economy, it is easier to identify changes in individual inflation forecasts due to a policy switch alone, without any additional effects due to group coordination dynamics. The disadvantage is that our experiment is silent

[1]The Riksbank was the only central bank in recent history to undertake PLT in the 1930s. See Berg and Jonung (1999) for a lucid description of PLT in Sweden.

on interesting questions regarding the aggregate volatility costs of a possible policy transition from IT to PLT.

After this introduction, Section 2 reviews the related studies. Section 3 offers a discussion of our simulation models leading to the main hypothesis statement. Section 4 describes the experimental design and procedures. Section 5 reports the results. Section 6 concludes.

2. Related Literature

In addition to the experiments investigating the stability of PLT mentioned earlier, our paper is closely related to other experimental studies analyzing inflation expectations in the New Keynesian environment. Several examples are Adam (2007), Luhan and Scharler (2014), Petersen (2014), Pfajfar and Žakelj (2014, 2018), Cornand and M'baye (2018a, 2018b), Hommes, Massaro, and Weber (2019), Kryvtsov and Petersen (2019), Mirdamadi and Petersen (2019), Assenza et al. (2021), Cornand and Hubert (2022), and Petersen and Rholes (2022). Most of these papers admit some level of interaction within groups of participants as aggregated individual forecasts affect simulated model dynamics.[2] Equilibrium stability issues and coordination patterns under various policy frameworks and policy communication treatments are often the main focus. As mentioned in the previous section, our paper focuses solely on individual forecasts in response to exogenously simulated rational expectations equilibrium.

At a broader level, our paper falls within the very active and fruitful "learning-to-forecast" area of experimental research. Some of the recent papers not mentioned above are Roos and Luhan (2008), Anufriev and Hommes (2012), Lambsdorff, Schubert, and Gamattei (2013), and Kryvtsov and Petersen (2021). Our analysis of individual-specific forecasting strategies in Section 5 builds on insights from these influential studies.

Recent surveys of experimental economics literature, such as Hommes (2011), Amano, Kryvtsov, and Petersen (2014), Cornand

[2]One exception is Pfajfar and Žakelj (2014), who test the rationality of inflation forecasts in an experiment without feedback effects of individual forecasts on the simulated model.

and Heinemann (2014, 2019), Duffy (2017), and Bao, Hommes, and Pei (2021), attest to a wide variety of questions from finance, macroeconomics, monetary economics, and central banking that are fruitfully investigated within a laboratory environment. There is a growing recognition among policymakers that experimental economics is a promising alternative for testing new policy tools, frameworks, and payment methods (Camera 2017). We hope our paper offers useful insights regarding the practicality of history-dependent monetary policies.

3. Simulation Models and the Main Hypothesis

Following most of the recent theoretical and experimental literature comparing monetary policy regimes, we use the canonical New Keynesian environment to model IT and PLT equilibria. As stated in the introduction, in our experiments, individual inflation expectations do not affect macroeconomic outcomes generated from a rational expectation equilibrium of the simulated model. The subjects of the experiments observe realized past values of macroeconomic variables in the model economy (in which all firms and households behave rationally) and attempt to predict the current inflation rate. The accuracy of subjects' inflation forecasts determines the amount of their payout. We aim to see if experimental participants adjust their expectations after an announced switch from IT to PLT, producing inflation forecasts consistent with the target-reverting nature of the price level under PLT.

We conducted two sets of experiments. The main experiment is a simplified forecasting environment in which either the inflation rate (IT) or the price level (PLT) follows an i.i.d. stochastic process. The second set of experiments, in which we consider a more realistic inflation forecasting problem, is a much richer forecasting environment in which inflation (and price-level) fluctuations are persistent, and other variables, such as output, interest rates, and nominal wages, are relevant for rational expectation (RE) inflation forecasts. The fully structural model, used in our experiments with a richer forecasting environment, features price-setting monopolistically competitive firms, a central bank, and utility-maximizing households. The central bank follows an interest rate rule, which reacts to inflation rate fluctuations (under IT) or price-level fluctuations (under PLT) around their respective time-invariant targets. The monopolistically competitive firms face time-dependent price-adjustment frictions, giving rise to Taylor-style staggered price setting.

Although our simpler forecasting environment uses i.i.d. stochastic processes, the generated inflation and price-level dynamics are consistent with the canonical New Keynesian model (derived from the fully articulated structural model) summarized by two structural equations,

$$\pi_t = \beta E_t \pi_{t+1} + \kappa x_t + u_t \tag{1}$$

$$x_t = E_t x_{t+1} - \gamma \left(R_t - E_t \pi_{t+1}\right) + g_t, \tag{2}$$

plus an additional relation characterizing monetary policy,

$$R_t = \delta^{IT} \left(E_t \pi_{t+1} - \pi^T\right) \text{ under IT} \tag{3a}$$

$$R_t = \delta^{PLT} \left(E_t p_{t+1} - p^T\right) \text{ under PLT,} \tag{3b}$$

where p_t is the log-price level and p^T is its fixed target. Likewise, $\pi_t = p_t - p_{t-1}$ is the inflation rate and π^T is the inflation target. The other variables are the output gap x_t, the cost-push shock u_t, the nominal policy rate R_t, and the demand shock g_t. Importantly, $E_t \pi_{t+1}$ is the model-consistent rational expectation of inflation, not to be confused with individual inflation forecasts of experimental subjects. The first structural equation is the New Keynesian Phillips curve relationship (i.e., the supply curve), while the second structural equation is the dynamic IS relation (the demand curve), which relates the output gap with the real interest rate $(R_t - E_t \pi_{t+1})$. The positive parameters β, κ, and γ are derived from the structural model parameters.

The relevant monetary policy depends on the policy framework and stabilizes inflation or price-level fluctuations around their respective targets. The inflation and price-level fluctuations might be larger or smaller depending on the values of the response coefficients, δ^{IT} and δ^{PLT}. Our assumption of i.i.d. inflation under IT (price level under PLT) requires a very high positive value of δ^{IT} (δ^{PLT}), implying an extremely aggressive response to expected future deviations of inflation (price level) away from the target. High positive values of response coefficients δ^{IT} (δ^{PLT}) cause fluctuations in $E_t \pi_{t+1} - \pi^T$ ($E_t p_{t+1} - p^T$) to be infinitesimal.

While such aggressive stabilization policies are unrealistic, the main advantage of assumed i.i.d. inflation and price-level processes is that they make rational expectations forecasts extremely simple. Under IT, the rational expectation forecast of the inflation rate is just the inflation target: $E_t \pi_{t+1} = \pi^T$, which for simplicity we set at zero, $\pi^T = 0$. Thus, knowing the inflation target is sufficient for making optimal (i.e., RE) inflation forecasts[3] under this simplest form of IT.

Likewise, under PLT, the rational expectation forecast of the price level is just its target, $E_t p_{t+1} = p^T$, which implies that the rational expectation forecast of inflation rate under PLT is the inflation rate that brings the price level back to target, $E_t \pi_{t+1} = p^T - p_t$. Thus, the subjects only need to know the current price level and the price-level target to infer the optimal (equivalently RE) inflation forecast under PLT.

The conceptual simplicity of optimal forecasts should simplify the forecasting problem considerably, provided the central bank's inflation (price-level) targeting objective is communicated sufficiently well.

The additional advantage of our i.i.d. assumptions is that they make the difference between optimal forecasts under IT and PLT extremely sharp. The most recent observed price level is relevant for optimal inflation forecasts under PLT and completely irrelevant under IT. In order to test if subjects understand the dynamics of inflation implied by PLT, we need to see if the price level is a significant determinant (with the correct sign) of individual inflation forecasts under PLT (treatment) but not under IT (control). Consequently, we can state our main test hypothesis as follows:

MAIN HYPOTHESIS: *The subjects of our experiments, who experience a sudden policy switch from IT to PLT, show a failure to understand the implications of the new PLT framework for inflation dynamics by not adjusting their forecasting strategy to exploit a negative correlation between the optimal forecast of inflation,* $E_t \pi_{t+1}^{PLT}$, *and the last observed price-level deviation from target,* $p_t^{PLT} - p^T$.

[3]In our experiments, rational expectations forecasts of inflation are optimal forecasts. Thus, we use these two terms interchangeably.

Notice that our hypothesis has a relatively low-level requirement for individual forecast rationality. To reject this hypothesis, we need to see that the price-level fluctuations from the model have a significant negative correlation with individual forecasts under PLT but not under IT. Given the weak general support for rational expectations in the empirical literature, we think our approach is prudently conservative.

As we will see in the results section, testing PLT understanding in the richer model environment is much less straightforward, as the optimal forecasts of persistent inflation fluctuations, with many more forecast-relevant variables, are much harder to learn. Nevertheless, we think a more realistic forecasting environment is important to enhance our study's external validity and policy relevance.

4. Experimental Design

The simple and the richer model experiments were conducted at the Center for Interuniversity Research and Analysis on Organizations (CIRANO) experimental economics laboratory. Subjects were recruited from a standard "convenience" subject pool consisting of university undergraduates using the widely used ORSEE Internet recruiting tool (see Greiner 2004 for details), which provides access to students at four universities in Montreal, Quebec, Canada. The experimental participants reported a wide variety of intended majors and were roughly equally split between genders. We view our subject pool as fairly representative of a general household population, the majority of which is not typically faced with price setting or professional forecasting activities. Thus, our experiments can be viewed as testing whether the general household population is capable of aligning their inflation expectations with PLT in an economy where (i) the policy switch from IT to PLT is well publicized and the PLT price-level target is known, and (ii) professional price setters in the financial markets and the product markets ensure inflation dynamics that are predominantly consistent with rational expectations.

4.1 Inflation Dynamics and Payoffs in the Simple Experiment

In our simple forecasting environment, we use a random number generator to generate normally distributed time series for an i.i.d.

inflation rate $\pi_t^{IT} \sim N\left(0, 0.8^2\right)$ under IT or an i.i.d. price level $p_t^{PLT} \sim N\left(0, \frac{0.8^2}{2}\right)$ under PLT. Both series imply the same standard deviation of the inflation rate of 0.8 percentage point[4] since $var\left(\pi_t^{PLT}\right) = var\left(p_t - p_{t-1}\right) = 2var\left(p_t\right) = 0.8^2$. We intentionally held the variability of inflation the same across IT and PLT regimes so that any gains in inflation forecasting under PLT would not be attributable to a change in inflation volatility.

As mentioned earlier, under IT, we assume that the central bank stabilizes expected inflation at zero, that is, $E_{t-1}\pi_t^{IT} = 0$. Under PLT, on the other hand, the central bank sets policy such that the expected price level is stabilized at its constant target, that is, $E_{t-1}p_t^{PLT} = p^T$.

There are three features of the current approach worth highlighting: (i) This experimental design matches the targeting horizon of the central bank and inflation forecasting period at one period such that the rational expectations inflation forecast is conceptually simple and does not require elaborate forecasting techniques to project the dynamic path of inflation multiple periods into the future; (ii) as the optimal forecasts are simple, any deviations from rational forecasts are straightforward to visualize and analyze; and (iii) since optimal forecasts of inflations have drastically different dynamics under IT and PLT (optimal inflation forecast is fixed under IT and it is perfectly negatively correlated with the current price level under PLT), the experiment allows us to clearly determine if participants exploit the additional information that the price level provides under PLT.

We discretized the forecasting strategy space to simplify and focus the subjects' decision-making on the essence of the forecasting problem. A participant is asked to choose an interval where he or she predicts the next period's inflation rate will lie. If the correct interval is chosen, the payoff is maximized. There are 13 intervals comprising 11 interior intervals and 2 unbounded intervals at the endpoints. The interior intervals span 0.5 percent, with the middle

[4]The standard deviation of 0.8 is slightly smaller than the 0.93 percentage point observed between January 1995 and December 2009 for the year-over-year Canadian CPI inflation rate.

interval centered on 0 percent inflation and bounded by 0.25 percent and –0.25 percent. The other interior intervals are constructed similarly. The two endpoint intervals capture extreme inflation forecasts greater than 2.75 percent or less than –2.75 percent. The discretized decision problem may introduce some approximation errors into our empirical estimation results, as ranges of inflation realizations within each interval are represented by single values: midpoints for interior intervals and plus or minus 3 percent for endpoint intervals. A half percentage point width of each interior interval offers a good compromise between simplicity and accuracy, as we found a high degree of correlation between the actual inflation realizations in our simulated model and their discretized approximations. This correlation was higher than 0.968 across our IT and PLT inflation series.

A quadratic loss function based on forecasting accuracy determines payoffs. Incorrect intervals result in increasingly smaller payoffs depending on the distance from the interval containing the correct ex post realization of inflation.[5] Note that the correct inflation realization is unknown ex ante, as it depends on random realizations of i.i.d. stochastic innovations. Nevertheless, since i.i.d. innovations are unpredictable and have a zero mean, it is always optimal to select the interval which contains the rational expectation forecast of inflation in the simulated model, $E_{t-1}\pi_t$.

Owing to the flatness of a quadratic function at its maximum, the quadratic payoff function used in this experiment rewards inflation forecasts that are "close" almost as well as inflation forecasts that lie in the correct interval. The latter feature suggests that it may be more difficult for subjects to detect small deviations from optimal inflation forecasts. On the other hand, large forecasting errors are harshly penalized, creating some risks of discouraging participants.[6]

[5]Every participant received at least the $10 participation payoff after the experiments were concluded, so negative total session earnings were eventually ignored.

[6]Lambsdorff, Schubert, and Giamattei (2013) discuss the relative merits of quadratic and linear loss functions in learning-to-forecast experiments. Our choice of a quadratic loss function is congruent with theoretical studies of monetary policy, which often use quadratic loss functions to approximate the central bank's objectives.

4.2 Experimental Procedures and Parameters

As a first step, we stochastically simulate the inflation and the price level under IT and PLT scenarios and store the resulting time-series data. In an IT regime, we simulate an economy where the central bank targets a zero inflation rate. In a PLT regime, the central bank stabilizes the price level at a constant price-level target.

We implement our simulated macroeconomy over the computer network in the CIRANO experimental laboratory. The computer provides an interface that presents subjects with the history of the previous eight periods of model-simulated inflation and the aggregate price level. These two variables are displayed in tables and graphs on the subjects' computer screens. The only other variables shown on the screens are (i) their individual inflation forecasts, (ii) their resulting period-by-period profits (i.e., payoffs) earned in the previous periods, and (iii) their progressive cumulative profits. No other variables, like output gap or interest rates, are displayed, as they are not necessary for making optimal (RE) forecasts. Indeed, by providing only inflation and the price level, we aim to focus the subjects' attention solely on variables relevant to optimal inflation forecasts.

At the beginning of each experiment, subjects are shown eight consecutive periods of inflation rate and price-level realizations and asked to predict inflation for the current period. In the case of IT, a horizontal line fixed at zero in the inflation graph reminded the subjects of the central bank's inflation target; in the case of PLT, a line showing the price-level target on the price-level graph was shown instead. After making their choice by selecting one of the 13 forecast intervals, the next period's results are displayed, always providing a window with the last eight periods' results.

Subjects were instructed that their task was to predict the inflation rate in a computer-simulated economy. To clarify the role of the central bank, we made the point that under IT, the central bank is not concerned with the past price level. Under PLT, subjects were instructed that the central bank would act to bring the price level to its constant target. The two regimes' instructions were parallel, involving identical paragraph and sentence structures.[7]

[7] The appendix contains the experimental instructions.

Finally, subjects are given a broad overview of the underlying model used to generate the time-series data but not its details. They were also told that random shocks present in the model would make it harder to discover the economy's structure. This type of instruction, where details of a complicated macro model are not revealed, has been used in many recent learning-to-forecast models and reflects that people in the real world make forecasts of inflation without a complete understanding of the macroeconomy.

We focused on two experimental treatments simulating a policy regime switch from IT to PLT: (i) the *manipulated* IT-PLT treatment and (i) the *control* IT-IT treatment. Both treatments were preceded by a 20-period *training* period, which allowed the subject to "practice" under an initial IT regime. For these first 20 periods, subjects repeatedly predicted inflation for the current period without pay. The absence of payoff implications allowed subjects to learn the system without fear of being penalized for experimentation.

The practice session was followed by two consecutive profit-incentivized segments, lasting 40 periods each. The first 40 periods maintained the same IT policy regime, but the sessions were twice as long as the practice session, and the subjects were now paid for the accuracy of their forecasts. Depending on the treatment, the second 40 periods were conducted under either an IT or a PLT regime.

Thus, after acquiring experience in a shorter training session designed to maximize experimentation and learning, subjects repeated the task with an economic incentive to make an accurate inflation forecast. More specifically, in the control (IT-IT) treatment, subjects made their forecasts in three consecutive inflation-targeting economies, whereas in the manipulated (IT-PLT) treatment, subjects practiced under inflation targeting, forecasted for pay under inflation targeting, and then forecasted for pay after an announced regime change to price-level targeting. Note that regardless of whether a regime change is implemented, the session is paused after the first segment, and a single page of instructions is distributed for the latter 40 periods. The instructions reminded subjects of the role of the central bank and stated that the central bank either continued with IT or shifted to a PLT regime. Both sets of instructions contained explicit guidance regarding the inflation targeting or the price-level targeting objective of the central bank. Given the recent findings in Honkapohja and Mitra (2020) and Coibion et al. (2023)

discussed in the introduction, such guidance is likely crucial for subjects' understanding of the difference between inflation dynamics under IT and PLT.

One potential disadvantage of revealing the central bank's objectives is that our experimental subjects might condition their forecasting behavior on their prior knowledge regarding the IT monetary policy framework prevailing in Canada. An alternative approach—followed, for example, by Engle-Warnick and Turdaliev (2010)—is to recast the decision-making problem into another seemingly unrelated activity to break any association with monetary policy. We did not follow this strategy because we were interested in seeing if the subjects experiencing an IT framework, both in the real world and in the lab, would adjust their expectations after a simulated switch to PLT. We see prior knowledge of IT as an important ingredient of our subjects' information set.

As mentioned, subjects are paid depending on the distance between their chosen interval and the correct realized interval. The payoff function is set using Monte Carlo simulations so that optimal forecasting would yield approximately \$40 on average for 80 forecasting periods. Twenty-nine subjects participated in the IT-IT treatments, and 25 participated in the IT-PLT targeting treatment. The average payoff was approximately \$30 in addition to a standard \$10 show-up fee at the experimental laboratory.

5. Results

5.1 Graphical Analysis

In this section, we use histograms to compare the distributions of participants' actual inflation (interval) predictions against optimal forecasts. Subjects' payoffs depend on deviations of actual inflation realizations in the model from their individual one-step-ahead forecasts. Due to shocks in our model, optimal forecasts, which are also rational expectations forecasts, do not always result in a maximum payoff. "Lucky" guesses could be suboptimal ex ante but ex post precise. Our focus is not on luck but on individuals' ability to form rational forecasts.

Figure 1 plots the histogram summarizing the distribution of deviations of chosen inflation intervals from their (ex ante) optimal

Figure 1. Deviations from the Rational Expectation Forecasts in the IT-PLT Session

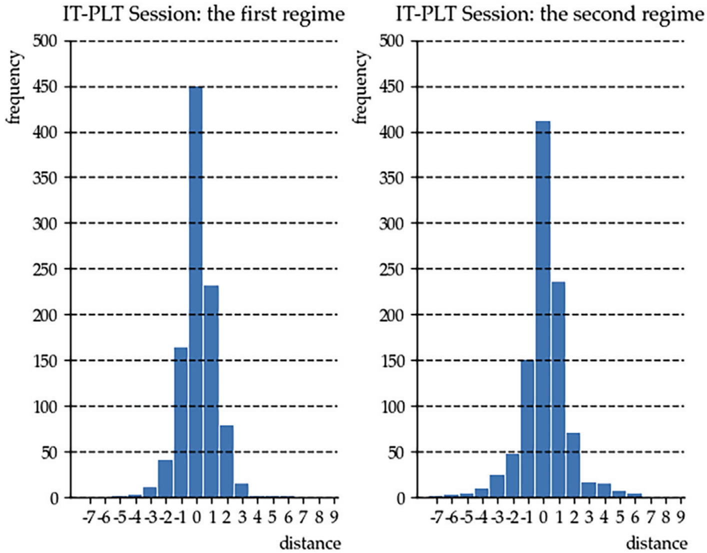

Note: The IT (PLT) is the first (second) regime. The horizontal axis shows the number of intervals between the chosen bin and the bin containing the optimal inflation forecast. The total number of deviations is 1,000.

(RE) counterparts in the IT-PLT session. The horizontal axis measures the integer distance between the chosen bin (predicted interval) and the optimal one (i.e., the interval containing the ex ante optimal inflation forecast). A positive (negative) distance means that the predicted bin contains higher (lower) inflation values than the bin containing the optimal inflation forecast. The tallest bars, labeled as zero, show the number of cases when subjects selected inflation intervals containing the optimal inflation forecast value. There were 450 (412) such cases out of 1,000 possible cases in the IT (PLT) regime.[8] The following bar to the right (left), labeled "1" ("–1"), shows the number of cases where chosen intervals were adjacent right above (below) the interval containing the optimal forecast. The other bars

[8]Each panel's total number of choices equals 1,000, as 25 subjects made 40 choices under the IT regime and 40 under the PLT regime.

were constructed similarly, showing larger deviations from optimal forecasts.

We see from Figure 1 that the IT scenario admits a slightly less dispersed histogram than PLT trials, meaning that subjects deviated less, on average, from optimal forecasts under IT. It is unsurprising because they had a 20-period training session to learn IT, while PLT was wholly unexpected and had to be learned in real time. Surprisingly, however, the average payoffs were higher in the PLT session. More specifically, under PLT, the subjects earned, on average, 79.8 percent of the optimal forecast payoffs, whereas under IT, subjects earned 69.7 percent on average. The range of realized earnings was also wider under PLT: payoffs under PLT ranged from a \$48.96 loss to a \$4.71 gain, both relative to the optimal forecast payoff. Meanwhile, the corresponding range for IT was between a \$32.48 loss and a \$0.12 gain. These results appear to suggest that the higher average payoffs under the PLT game arose due to more active forecast experimentation with bigger losses and gains.[9] Notice that due to the randomness of inflation realizations, some lucky participants managed to earn more than possible if they had always formed rational expectation forecasts.

Figure 2 plots the histogram summarizing the distribution of deviations of chosen inflation intervals from their optimal counterparts in the IT-IT session. The histogram was constructed in the same way as Figure 1. We see from Figure 2 that the second IT segment (right panel) has a much more concentrated distribution with a substantially higher number of zero deviations from optimum than the first IT segment. The improved forecasting performance during the second IT period is also reflected in a higher average cumulative payoff of 82 percent (versus 67.2 percent in the first IT game) relative to the optimal forecasting strategy. Thus, some learning of the IT regime is apparent in the results of the IT-IT session. In fact, one subject followed the optimal forecast strategy of predicting zero inflation throughout the second IT trial.

It is an interesting question (although somewhat outside of our main focus on modeling the transition from IT to PLT) whether similar learning gains would be realized under PLT if subjects had

[9]The mean squared error (MSE) calculations give similar results. The session average MSE is lower under PLT than under IT, but the range of individual-specific MSE numbers is substantially wider under PLT.

Figure 2. Deviations from the Rational Expectation Forecasts in the IT-IT Session

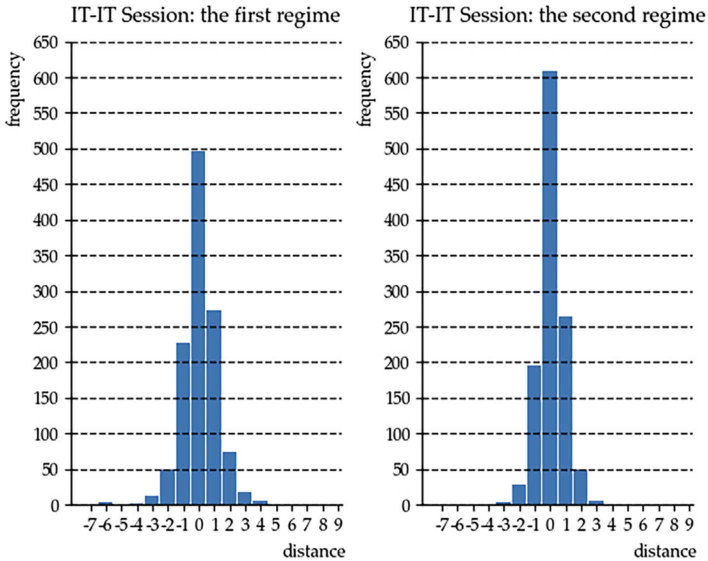

Note: The horizontal axis shows the number of intervals between the chosen bin and the bin containing the optimal inflation forecast of zero. The total number of deviations is 1,160.

more experience with forecasting inflation under PLT. Figure 3 shows the distributions of forecast deviations from optimal forecasts in a PLT-PLT session, preceded by a 20-period practice session without pay conducted under PLT. Both PLT games for pay lasted 40 periods each and followed a similar setup to an IT-IT session, with only the monetary policy framework being PLT throughout. Thirty-one subjects participated in the PLT-PLT session, giving 1,240 (31 times 40) total deviations in each of the two PLT panels shown in Figure 3. Once again, the right panel shows a much more concentrated distribution with a more pronounced peak at zero, suggesting improved forecasting performance in the second PLT portion of the experiment.[10]

[10]The differences in average cumulative payoffs and average MSE values support this conclusion that the second PLT part of the PLT-PLT session showed improved forecasting performance relative to the first segment.

Figure 3. Deviations from the Rational Expectation Forecasts in the PLT-PLT Session

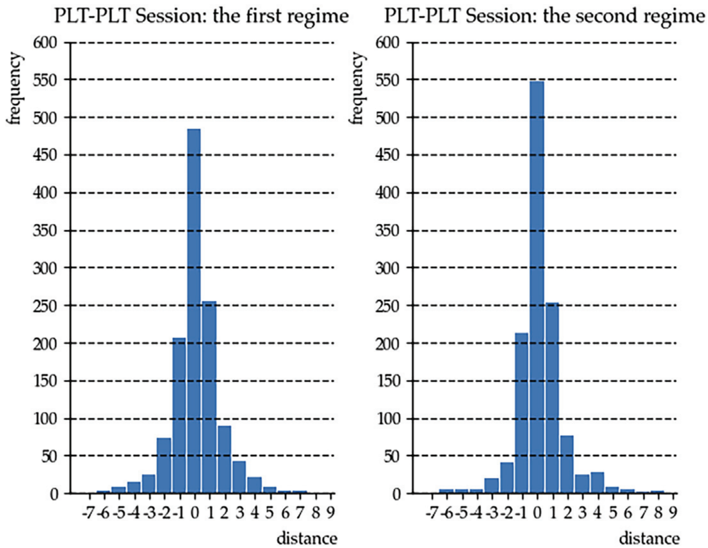

Note: The horizontal axis shows the number of intervals between the chosen bin and the bin containing the optimal inflation forecast. The total number of deviations is 1,240.

Interestingly, all six histograms shown in Figures 1, 2, and 3 have more observations to the right of zero, suggesting that people were more likely to overpredict inflation relative to optimal forecasts. This is likely the reflection of people's general awareness that the actual inflation in Canada is typically positive.

Returning to IT-PLT session results, Figure 4 compares the optimal PLT forecasts in the PLT game with the subjects' predictions without taking a difference between them. It highlights how subjects' inflation forecasts under PLT differ from their optimal counterparts. Specifically, individual inflation forecasts under PLT (left panel) are too concentrated relative to optimal forecasts (right panel). Participants choose the zero bin too often: 41.7 percent versus 32.8 percent under the optimal prediction scenario, suggesting they are not fully exploiting the implications of a mean-reverting price level for optimal inflation forecasts.

Figure 4. PLT Game from IT-PLT Session: Distribution of Individual Predictions (left panel) and Optimal Forecasts (right panel)

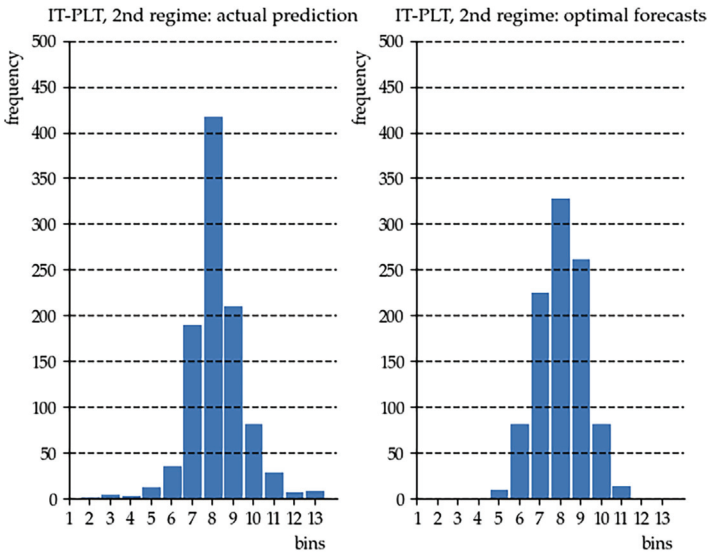

Note: The total number of observations is 1,000.

This overconcentration of actual forecasts under PLT contrasts with excessive forecast dispersion under IT, where optimal forecasts are always zero, but the actual predictions are dispersed. While it is unclear why the forecasts are overconcentrated in one case and too dispersed in another, it seems plausible that when faced with shocks and uncertainty, subjects put only a partial weight on information contained in inflation or price-level targets. If subjects put partial weight on guidance regarding central bank objectives while retaining some weight on other forecasting heuristics—like trend extrapolation, for example—then we should see IT forecasts dispersed relative to target. We might also see PLT forecasts that are too concentrated relative to optimal forecasts, depending on an alternative forecasting heuristic used.

We explore individual forecasting heuristics in the next section, in which we use subject-specific regressions to analyze individual prediction strategies.

5.2 Individual Forecasting Heuristics

In this section we follow insights from Anufriev and Hommes (2012), as well as from other studies, such as Roos and Luhan (2008), Lambsdorff, Schubert, and Giamattei (2013), Pfajfar and Žakelj (2014, 2018), Cornand and M'baye (2018a), and Assenza et al. (2021), to fit an array of simple learning and forecasting techniques to our subject-specific prediction data. The objective is to find the best-fitting expectation formation model, thus gaining insights into dominant forecasting techniques and identifying any changes in these techniques in response to a policy switch from IT to PLT. Specifically, we estimate the following rules, as specified in Cornand and M'baye (2018a) and Assenza et al. (2021), but adapted to our simpler forecasting environment with exogenous simulation dynamics, which do not depend on individual forecasts.[11]

- An anchoring and adjustment expectation formation rule admits some weight on the rational expectation forecast (anchor) of current inflation, $E_{t-1}\pi_t$, inferred from the guidance regarding the central bank's objectives. In our experiments, subjects' individual forecasts of current inflation, $\pi^i_{t|t-1}$, are based on information observed up to the previous period and include last-period observations of price-level deviation from the target $\left(p_{t-1} - p^T\right)$ and inflation π_{t-1}. The rational expectation forecast of current inflation is $E_{t-1}\pi^{PLT}_t = -\left(p_{t-1} - p^T\right)$ under PLT and just zero under IT, $E_{t-1}\pi^{IT}_t = 0$. The rest of the weight is placed on the last observed inflation realization, i.e., $\pi^i_{t|t-1} = w_1 E_{t-1}\pi_t + (1 - w_1)\,\pi_{t-1}$. For estimation purposes, we transform this rule into a less restricted version,

$$\pi^i_{t|t-1} - E_{t-1}\pi_t = \alpha_0 + \alpha_1\left(\pi_{t-1} - E_{t-1}\pi_t\right), \qquad \text{(M1)}$$

[11]In most learning-to-forecast studies, subjects need to make two-step-ahead forecasts because the current realized aggregate variables depend on current individual forecasts. The one-step-ahead forecasting problem in our study simplifies the learning rules.

in which we drop the requirement that α_0 has to be equal to zero, as the absence of an intercept could complicate the interpretation of our estimation results.

- The naive expectation formation rule uses only last-period inflation to form the current inflation expectation, irrespective of policy regime.

$$\pi^i_{t|t-1} = \beta_0 + \beta_1 \pi_{t-1} \tag{M2}$$

For the IT regime where $E_{t-1}\pi_t = 0$, the naive rule is equivalent to the above-mentioned anchoring and adjustment rule.

- The AR(1) expectation rule is a self-referential first-order autoregressive model.

$$\pi^i_{t|t-1} = \gamma_0 + \gamma_1 \pi^i_{t-1|t-2} \tag{M3}$$

- The trend-extrapolation model.

$$\pi^i_{t|t-1} - \pi_{t-1} = \delta_0 + \delta_1 \left(\pi_{t-1} - \pi_{t-2}\right) \tag{M4}$$

- The adaptive expectations model, $\pi^i_{t|t-1} = \pi^i_{t-1|t-2} + \sigma_1 \left(\pi_{t-1} - \pi^i_{t-1|t-2}\right)$. For estimation purposes, we transformed it into a less restrictive version,

$$\pi^i_{t|t-1} - \pi^i_{t-1|t-2} = \sigma_0 + \sigma_1 \left(\pi_{t-1} - \pi^i_{t-1|t-2}\right). \tag{M5}$$

Notice that all of our five rules, M1–M5, have a similar empirical formulation with an intercept and one slope coefficient, so we can use unadjusted R-squared to compare the goodness of fit across these models. Table 1 summarizes our findings regarding the best-fitting model for each game and each session.

One immediate observation from Table 1 is that the trend-extrapolation model (M4) is the dominant strategy across all sessions and both IT and PLT regimes, with the only exception being the first IT game (IT1) in the IT-IT session for which adaptive learning had the highest R-squared value for 16 people out of 29 participants. The AR(1) model was the best-fitting model for at most one person per session and was associated with the worst payoffs in two cases and mediocre results in the remaining two cases. The naive model

Table 1. Best-Fitting Expectation Formation Model: Simple Experiment

Session and Policy Regime	Number of Subjects with the Best Fit				
	Anchor M1	**Naïve M2**	**AR(1) M3**	**Trend M4**	**Adaptive M5**
IT-PLT:					
IT	5	0	1	10	9
PLT	6	1	0	14	4
IT-IT:					
IT1	3	0	1	9	16
IT2	4	0	0	15	10
PLT-PLT:					
PLT1	6	2	1	19	3
PLT2	7	0	1	23	0

(M2) was the best-fitting model for three people across two PLT games,[12] with relatively good payoff outcomes.

The adaptive model (M5) was the first or second most frequently used model under the IT policy regime but was dominated by the anchoring and adjustment model (M1) when the policy became PLT. In PLT-PLT sessions, for example, only three out of 31 subjects used the adaptive model in the first segment (PLT1), and none used the adaptive model in the second portion (PLT2).

Table 2 focuses on the IT-PLT session and shows the changes in the best-fitting model between the first IT regime and the second PLT regime. From the M1 column, we see that six people switched their forecasting strategy to anchoring and adjustment (M1), four from trend extrapolation (M4), and two from adaptive strategy (M5). From the M1 row of the M4 column, we see that five people changed their forecasting strategy from M1 (or perhaps from the observationally equivalent M2) followed under the IT policy regime to the trend-extrapolation model (M4) followed under PLT. Our

[12] Under IT, the naive model was observationally equivalent to the anchoring and adjustment model (M1), so the absence of M2 observations in the IT regime was by assumption.

Table 2. Expectation Formation Transitions: Simple Experiment, IT-PLT Session

Best-Fit Model, IT	Best-Fit Model, PLT				
	Anchor M1	Naïve M2	AR(1) M3	Trend M4	Adaptive M5
M1	0	0	0	5	0
M3	0	0	0	0	1
M4	4	0	0	5	1
M5	2	1	0	4	2

finding that people change their forecasting strategy after a monetary framework change is consistent with the findings of Bao et al. (2012). These authors find that an evolutionary learning model with endogenous changes in forecasting strategies fits well with the data from experiments in which subjects experience a large shock. A monetary policy framework switch in our model can be considered a big shock that changes the optimal forecasting strategy.

Overall, the information in Tables 1 and 2 suggests that the anchoring and adjustment model (M1) was used more frequently under PLT, but trend extrapolation (M4) remained the most popular choice. Nevertheless, even if the M1 model is not the best fit for most people, it could still be that most subjects understood that the latest observed price level is relevant (with a negative coefficient) for inflation forecasts under PLT and irrelevant under IT. In order to test this conjecture, we run individual-specific regressions of the following form:

$$\pi^i_{t|t-1} = q^i_0 + q^i_1 p_{t-1} + q^i_2 \pi_{t-1} + \varepsilon^i_t, \tag{4}$$

where $i = 1, 2, ..., 25$ for all of the participants in the IT-PLT session. Per our stated main hypothesis, an individual subject would show an inability to understand the implications of the PLT framework for inflation dynamics if the estimated $\hat{q}^i_1$ coefficient on the last observed price level p_{t-1} was not significantly different from zero or had the wrong (positive) sign. If the $\hat{q}^i_1$ estimate is negative and significant, it would also be helpful for robust identification if $\hat{q}^i_1$ was not significantly different from zero under IT, in sharp contrast to its significance under PLT. The lag of inflation is included to allow for the possibility that alternative forecasting strategies relying on the last observed inflation rate were also deemed relevant for individual predictions.

Our estimation results for the PLT game of the IT-PLT session indicates that out of 25 people who experienced a policy switch from IT to PLT, 22 subjects had a negative estimate of q_1^i, 18 of them were significant at the 5 percent level of significance, and 13 were significant at the 1 percent level of significance.

In the case of the IT game (of the same IT-PLT session), out of 25 people, 11 had negative estimates of q_1^i, of which three were significant at the 5 percent level of significance, and only one estimate was significant at the 1 percent level.

Thus, the individual-specific regression estimates suggest that most subjects in the IT-PLT session (at least 18 if we use the 5 percent level of significance) exhibited some understanding of the PLT framework for inflation dynamics despite having to learn the new policy framework after a sudden transition from IT.

5.3 Panel Estimation Results

In this section, we use panel estimation techniques to contrast the aggregate forecasting performance of our subjects in the manipulated IT-PLT session relative to that in the control IT-IT session. Thus, we attempt to infer at a more aggregated level whether inflation expectations adjust in a manner consistent with PLT by comparing panel estimation results from the second part of the IT-PLT session with those from the second segment of the IT-IT session. We allow the intercept to differ across subjects to control for individual forecasting heterogeneity according to the random effects panel estimation model. The random effects model is preferred to the fixed effects model, provided no cross-sectional correlation exists between the explanatory variables (regressors) and the individual random effects. Since our regressors are exogenous random draws from the same simulation model unaffected by individual forecasts, there is no reason to expect such a correlation in our results.[13]

[13]Regardless, the fixed effects model gives very similar estimation results.

Our (least restricted) panel regression is specified to have the same two regressors as the individual-specific regressions mentioned earlier:

$$\pi^i_{t|t-1} = \alpha^i + \beta p_{t-1} + \gamma \pi_{t-1} + \varepsilon^i_t. \tag{5}$$

However, it imposes the restriction that β and γ coefficients have common values for all individuals in the panel. The parameters β and γ measure the sensitivity of inflation predictions to movements in the price level and inflation, respectively. Under IT, we expect to see $\hat{\beta}$ and $\hat{\gamma}$ equal to zero if expectations are rationally generated since optimal inflation forecasts in the simple IT framework are always zero. Under PLT and rational expectations, we would expect $\hat{\beta} = -1$ and $\hat{\gamma} = 0$, as the best prediction of inflation is perfectly negatively correlated with today's deviation of the price level from its target. Further, to account for potential heteroskedasticity and serial correlation of individual forecasting errors, we use heteroskedasticity and serial correlation robust standard errors introduced by Newey and West (1987).[14]

The estimation results are reported in Table 3. Columns 1 and 3 provide a sense of effects of p_t and π_t on predicted inflation under IT. The estimates of γ, capturing the sensitivity to past inflation, are approximately equal to 0.145 and are highly statistically significant in both regressions, implying that subjects tend to forecast inflation as if it were persistent even though it is not.[15] The overall level of inflation prediction errors is relatively high, with incorrect intervals chosen in 55 percent of the cases (see Figure 1, left panel). Interestingly, we see similar perceived inflation inertia under PLT (column 4 of Table 3). The elasticity of expected inflation with respect to the

[14] Prediction errors for a minority of subjects were serially correlated with significant AR(1) coefficients. Interestingly, fewer people exhibited serial correlation of prediction errors in the IT regimes than in the PLT regimes.

[15] This phenomenon of observing persistence in random data has been documented in the psychology literature and is referred to as the "hot hand" (see, for example, Tversky and Kahneman 1974). The experimental economics literature has also noted this behavior. Huber, Kirchler, and Stockl (2010) developed a laboratory experiment where subjects attempted to predict an unknown process (in fact, an i.i.d. process). After this experience, subjects advised a second group of experimental participants on predicting the same unknown process. Consistent with our results, Huber, Kirchler, and Stockl (2010) observe advice corresponding to predicting a persistent series.

Table 3. Random Effects Panel Regression Results for Inflation Predictions: IT-PLT versus IT-IT

Independent Variables	Policy Regime (1) IT	(2) PLT	(3) IT	(4) PLT
Price Level		–0.375* (0.041)	–0.002 (0.005)	–0.554* (0.041)
Inflation	0.145* (0.016)		0.146* (0.016)	0.175* (0.034)
R-squared	0.073	0.104	0.073	0.127
Number of Observations	1,160	1,000	1,160	1,000

Note: * indicates significance at the 1 percent level. All regressions contain an unreported constant term. Standard errors in parentheses are based on Newey and West (1987).

latest realized inflation is about 0.175 (after controlling for the price level). It is not statistically different from the estimated γ coefficients in columns 1 and 3. The finding that past inflation is used to forecast future inflation was also noted by Cornand and Hubert (2020), who find this to be a common feature not only in experiments but also in surveys of inflation expectations of households, industry forecasters, professional forecasters, financial market participants, and central bankers. The authors interpret this commonality as evidence of the external validity of experimental inflation forecasts.

The effect of price-level deviations and the inflation rate on predicted inflation under PLT is found in columns 2 and 4. Recall that a weakly rational, as defined in our main hypothesis, forecast of inflation under PLT should, at least, exhibit a negative correlation with p_t. The results suggest that subjects use this directionality property of PLT, on average, by admitting negative estimates of β ranging between –0.37 and –0.55. These regression coefficients suggest that the inflation forecasts are suboptimal, on average, since both estimates of β are statistically different from –1. In other words, subjects are not taking full advantage of a key property of PLT in forming their inflation forecasts. From our individual-specific regressions discussed earlier, we know that several subjects failed to exhibit evidence of understanding the implications of PLT for inflation dynamics. Overall, subjects seem to produce inflation forecasts consistent with the stationarity implication of PLT. However, the accuracy of those forecasts does not improve over IT, as deviations from optimal forecasts occur approximately 59 percent of the time (see Figure 1, right panel).

5.4 The Richer Experiment

In an effort to enhance the external validity and policy relevance of our study, in this section, we report panel regression results based on a richer, more realistic experiment. There were 55 subjects in the IT-IT control treatment and 58 in the IT-PLT manipulated treatment. Many of the implementation details were similar to the simpler experiment. Subjects still had a chance to practice for 20 periods without pay under an IT regime and then forecasted the one-step-ahead inflation for money in a 40-period IT game. After this portion of the experiment was completed, a shift in monetary policy to PLT was announced and subjects were asked to continue forecasting one-step-ahead inflation. The control group did not experience a monetary policy change and completed an IT-IT session. A quadratic loss function, similar to that of the simpler experiment, calculates payoffs. The loss function was calibrated to yield approximately $20 per 40-period game under a rational expectations forecasting strategy.

We used a calibrated structural macroeconomic model based on Galí and Gertler (1999) to generate inflation, output, wage, and other series for our experimental study. Since the model is quite standard in the recent monetary policy literature, we are brief in describing its main features.[16] There are four types of agents in the economy: (i) the representative household that maximizes lifetime expected utility by choosing consumption, labor hours, and risk-free nominal bond holdings; (ii) perfectly competitive final good firms which produce final consumption goods from various intermediate good inputs using a constant elasticity of substitution (CES) production function; (iii) monopolistically competitive intermediate

[16]Detailed information on the macroeconomic model and its calibration can be found in Appendix A of the working paper version of this article, available at http://www.bankofcanada.ca/wp-content/uploads/2011/09/wp2011-18.pdf.

good producers producing different varieties of intermediate goods by utilizing labor hours supplied by the households in a diminishing returns to labor production technology, and (iv) a monetary authority that sets the nominal interest rate via an interest rate rule. Depending on the monetary policy framework in place, the monetary policy rule responds to inflation or price-level fluctuations relative to their targets. Importantly, for the external validity of our study, the average inflation in our model is 2 percent per year under both IT and PLT regimes, as targeted by many central banks, including the Bank of Canada. In addition to these price stability objectives, the monetary policy rule responds to output gap fluctuations. Producers of the intermediate goods are assumed to set prices according to Taylor-style staggered nominal contracts of fixed duration, yielding nominal rigidities in the model and giving a welfare-relevant stabilization role to monetary policy. The model abstracts from capital accumulation and fiat money.

Relative to the earlier setup, the current experimental design does not force the next period's inflation or price level back to its target within one period. Instead, the path of inflation under IT is dictated by the calibrated historical monetary (Taylor) rule. Under PLT, a similar rule is chosen and calibrated to maintain the same variance of inflation as under IT. The gradual adjustment of inflation or the price level back to target requires subjects to face the more difficult task of accounting for the dynamic path of inflation when calculating their inflation expectation. Adding to the difficulty, this experiment also requires subjects to provide point inflation forecasts rather than choose intervals, and the targeted price-level path grows at 2 percent rather than being a fixed value. Moreover, many more relevant variables had to be considered to form a rational expectations forecast. As per instructions:

> "To maximize your earnings, you will need to pay attention to the following variables, which will be displayed in their own graphs on the screen. All of these variables provide the same information to you as they do in the real economy.
>
> (1) Inflation—computed on an annual basis (over four quarters)
>
> (2) Output—the total real amount of goods and services produced by the economy. If output is above 100, then the economy

is producing more output than normal, and if it is below 100 it is producing less than normal.

(3) Price level—the average price of goods in the economy

(4) Interest rate—the short-term interest rate set by the central bank

(5) Wage—the average wage of the agents in the economy"

Additional instructions provided a sense of relationships between the variables:

"The underlying model of the economy determines the relationships between all of these variables, along with some randomness that makes the relationships somewhat difficult to determine. You will not be told these precise relationships. You will learn them through the simulation. Roughly speaking, the variables have these relationships to inflation: (1) inflation itself has some persistence, or inertia, which limits how much it can change from quarter to quarter; (2) output above 100 can signal higher inflation, and below 100 can signal lower inflation; (3) increasing price levels indicate inflation; (4) the central bank raises its interest rate to eventually bring down inflation and lowers it to eventually bring up inflation; (5) higher wages can be an indication of higher inflation."

In short, the current forecasting problem is more difficult than the previous case, imitating, to some extent, real-world complexities. Expectedly, individual forecasting performance varied considerably, but nevertheless, some experimental subjects produced surprisingly good forecasts. For example, the best-performing forecaster across the IT sessions earned 87 percent of the optimal strategy payoff. Similarly, the subject with the best forecasts across the PLT sessions scored about 88 percent of the optimal payout. Both examples illustrate impressive forecasting performance and suggest that a relatively high persistence of inflation helped, to some extent, make trend-extrapolation techniques successful at predicting the inflation cycles.

Once again, we summarize the data from the richer experiment using random effects panel regressions. The least restricted panel

Table 4. Random Effects Panel Regression Results for Inflation Predictions: IT-PLT versus IT-IT

Independent Variables	**Policy Regime**			
	(1) IT	**(2) PLT**	**(3) IT**	**(4) PLT**
Price Level		11.840* (1.177)	–0.044 (0.077)	–0.684 (0.545)
Inflation	0.749* (0.012)		0.749* (0.012)	0.747* (0.014)
R-squared	0.731	0.101	0.731	0.751
Number of Observations	2,145	2,262	2,145	2,262

Note: * indicates significance at the 1 percent level. All regressions contain an unreported constant term. Standard errors in parentheses are based on Newey and West (1987). The number of observations is 39 times the number of subjects.

regression is specified to have a similar form to that in the simpler experiment,

$$\pi^i_{t|t-1} = \alpha^i + \beta \left(p_{t-1} - p^T_{t-1}\right) + \gamma \pi_{t-1} + \varepsilon^i_t, \tag{6}$$

except the price-level target p^T_{t-1} is now growing over time. Thus, it needs to be subtracted from the actual price level to calculate the price-level deviation from its moving target. The parameters β and γ measure the sensitivity of inflation predictions to movements in the price-level deviation and inflation, respectively. Under IT, we expect $\hat{\beta}$ to be close to zero, as the price level is irrelevant for expected inflation under inflation targeting. Under PLT and rational expectations, we expect $\hat{\beta}$ to be negative to reflect the negative correlation between the expected inflation and the current price-level gap relative to its target. It is unclear what the estimated value of $\hat{\gamma}$ should be in either regime, but to the extent there is persistence in inflation dynamics, the likely value of $\hat{\gamma}$ is positive.

Table 4 shows the estimation results comparing the second IT regime from the IT-IT session (the control) with the PLT regime from the IT-PLT session (the treatment). The estimation results

Table 5. Random Effects Panel Regression Results for Optimal Inflation Predictions: IT-PLT versus IT-IT

	Policy Regime			
Independent Variables	**(1) IT**	**(2) PLT**	**(3) IT**	**(4) PLT**
Price Level		9.013* (1.318)	–0.151 (0.251)	–2.967* (0.658)
Inflation	0.701* (0.012)		0.701* (0.012)	0.781* (0.012)
R-squared	0.562	0.051	0.562	0.639
Number of Observations	2,145	2,262	2,145	2,262

Note: * indicates significance at the 1 percent level. All regressions contain an unreported constant term. Standard errors in parentheses are based on Newey and West (1987).

offer somewhat ambiguous conclusions. The estimates of $\hat{\gamma}$ are consistent and highly significant across columns 1, 3, and 4, suggesting that subjects put a stable weight on past inflation outcomes when formulating their current inflation predictions. The estimates of $\hat{\beta}$ are not consistent across columns 2 and 4: the estimate in column 2 has the wrong sign, while the estimate in column 4 has the correct negative sign but is not significant.

One way to check if the estimated coefficients make sense is to estimate the same regression on the model-generated rational expectation forecasts of inflation. Table 5 shows the estimation results for this benchmark case.[17] The general pattern of estimated coefficients is similar across Tables 4 and 5, except the price-level coefficient $\hat{\beta}$ in column 4 of Table 5 is much higher in absolute value and highly significant.

[17] Since no cross-sectional heterogeneity is present in model-simulated forecast data, the random effects model gives identical estimates as a simple pooled ordinary least squares (OLS) regression in columns 1, 3, and 4 of Table 5. In column 2, the random dffects coefficients differ from the pooled OLS estimates, likely because of the omitted variable misspecification in this regression.

One has to be careful about the interpretation of estimates in Table 5 because these are just reduced-form regressions, which have no clear theoretical interpretation. Under PLT, the structural model has five state variables—two endogenous and three exogenous shocks. The lag of the price-level deviation from the target, $\left(p_{t-1} - p^T_{t-1}\right)$, is one of the endogenous state variables under PLT, but the lag of the inflation rate π_{t-1} is not. For this reason, the reduced-form regressions estimated in columns 2 and 4 of Table 5 are likely misspecified. Nevertheless, it is somewhat reassuring that the signs of estimated $\hat{\beta}$ coefficients are consistent across Tables 4 and 5.

Overall, the results from the richer model do not provide conclusive evidence that subjects understand the directionality implication of PLT for expected inflation dynamics, primarily because the estimate of $\hat{\beta}$ in column 4 of Table 4 is not significant. This lack of precision and general ambiguity of results from the richer model highlights again the need for sharper differences between IT and PLT frameworks in future experimental studies. Our simpler model formulation was a step in that direction.

6. Concluding Remarks

The theoretical advantages of price-level targeting in stabilizing economic activity near the zero lower bound on interest rates generated substantial attention to this promising policy framework among academics and U.S. monetary policy officials. However, the new evidence presented in Coibion et al. (2023) regarding the limited understanding of average inflation targeting, the new policy framework adopted recently by the U.S. Federal Reserve Bank, cast doubts on whether the general public would understand the implications of PLT for inflation dynamics. The economic stabilization benefits of AIT and PLT frameworks rely crucially on economic agents' ability to understand their implications for inflation dynamics and adjusting their inflation expectations accordingly. Without such adjustments in expectations, AIT and PLT are likely to be destabilizing and welfare-reducing relative to the commonly practiced inflation-targeting framework. While the evidence presented in Coibion et al. (2023) is not very encouraging, the limited understanding of AIT could be the result of somewhat vague and unclear communication of the Fed regarding key implementation details of the new framework. However, PLT may be easier to visualize and communicate than AIT.

Given the potential stabilization benefits of PLT in a low interest rate environment, it seems useful to explore whether people would understand its implications for inflation dynamics after a transition from a monetary policy framework targeting the inflation rate. In the current paper, we attempted to shed some light on this issue by simulating a sudden transition from inflation targeting to price-level targeting in a controlled laboratory environment.

Our results suggest that inflation forecasting behavior changes after a sudden transition from IT to PLT as people learn to take advantage of the additional information contained in price-level fluctuations for expected future inflation dynamics. In our simple experiment, subjects move from (incorrectly) relying only on inflation to predict future inflation under IT to using (but not fully) the directionality implication of PLT under a PLT regime. Our investigation of individual forecasting strategies under PLT suggests that people attach some weight to the information contained in the current price level. However, they also often rely on other learning heuristics like trend-extrapolation techniques. Our experiment based on a richer, potentially more realistic, economic model was less conclusive. It still suggests that subjects tend to rely on the target-reverting nature of the price level to generate their inflation forecasts under PLT, but the regression coefficient on price-level fluctuations was not significant. It should be emphasized that the shift to PLT in the experiments was explained only once to subjects. In the real world, a central bank would likely undertake an ongoing communication strategy to explain and remind the public about the implications of PLT, thereby helping agents to exploit the properties of a PLT regime more fully. Moreover, real-world agents would also have access to inflation forecasts from professional forecasters. Incorporating these features into a macroeconomic experiment may be a useful avenue for future research. Another potential application of this framework is to explore whether the same decision-making problem recast as an exchange rate stabilization objective, perhaps moving from a floating exchange rate arrangement to a peg, would generate different outcomes because people may form their exchange rate expectations quite differently from the inflation rate expectations.

Finally, it is worthwhile to evaluate if a dual mandate of monetary policy, concerned with both price stability and output stability, could change the credibility of PLT.

Appendix. Instructions for the Simple Experiments

What You Will Be Doing

You will predict inflation in a computer-simulated economy. Your pay will depend on how accurate your predictions are. Your pay will depend only on your decisions and the results generated by the simulated economy. It will not depend on the decisions of any other participants. You make many decisions today. Each time you make a decision is called a period.

The Economy

You can think of the computer economy as simulating the activity of a real economy. It can be thought of as consisting of households, which work and buy goods; intermediate firms, which provide the goods needed to make a final good; the firm that produces the final good that is purchased by the consumers; and a central bank that uses its short-term interest rate to achieve control over the economy.

The Central Bank

The central bank provides stability to the economy. The bank has one mandate: to stabilize inflation. Inflation is the change of the average price of goods and services in the economy. It is the difference in the price level between the current period and the previous period. The central bank attempts to stabilize inflation at a target of 0 per period. That is, the central bank attempts to make the difference between the current price level and the previous price level 0. If inflation moves higher or lower than 0, due to randomness in the economy, the central bank will act to return inflation back to the target. The bank uses its interest rate to achieve its objective of stabilizing inflation at 0. This implies that the central bank is not concerned with achieving a price-level target but instead attempts to maintain or return inflation to its target. There will be a line on the screen showing you the inflation target.

How You Make Your Decisions

You will be shown the price level and inflation in your economy. The price level and inflation are determined by the structure of the economy and some randomness that makes the structure difficult to determine. At the start, the computer will show you eight periods of economic results. You will then predict inflation for the next period. After you do, the computer will show you results for the next period, and you will predict inflation again for the subsequent period. For your decision you choose an interval within which you expect the next period's inflation to fall. For example, one interval is 1.75–2.25 percent. In total there are 13 intervals. All of the intervals are 0.5 percent wide. You make your decision by clicking on the interval you choose. When you choose your bin, you will see an asterisk, that is, the character "*", located underneath the center of the bin containing the previous period's inflation. This character is placed on the screen to assist you with your forecast of inflation for the next period.

The Number of Decisions You Make

You will predict inflation for 20 periods for practice. You will not be paid for your practice. You may use the practice to learn how the simulation works.

You will then predict inflation for 80 periods for pay. For each decision the relationships between economic variables are identical. The randomness, however, will be different, and independent of any past decisions you make.

How You Will Be Paid

The better your prediction of inflation, the higher is your pay. Each period the computer will determine how many bins there are between your prediction and the bin that actual inflation falls in. The closer your predicted bin is to actual inflation, the higher your pay. This graph[18] shows your pay, in dollars, for a period, depending on how

[18]The graph of the quadratic payoff function is available in Appendix A of the working paper version. See http://www.bankofcanada.ca/wp-content/uploads/2011/09/wp2011-18.pdf.

many bins away your prediction is. The horizontal axis, labeled 0,1,2,. . .,7, represents the distance between the bin you choose and the bin within which inflation actually falls. The vertical axis, labeled 2.00, 1.00, 0.00,. . . ,–6.00 represents your pay for a decision in one period. For example, if inflation falls within the bin you choose, then the distance between your prediction and actual inflation is 0, and you earn $0.83 for the period. For another example, if inflation falls seven bins away from your prediction, then your pay for the period is approximately –$5.24. Since there are 13 bins, you could be as many as 12 bins off with your prediction, in which case you would earn approximately –$17.01 for the period. The pay is scaled so that on average, if you make the best possible prediction every quarter, you can earn about $40. During a period, it is possible to make negative earnings. Your earnings are computed by adding up your earnings for every period. You cannot make negative earnings in a session. You will be paid in cash for all of your decisions. The bottom line is that the better your prediction, the higher your pay.

A Reminder of the Role of the Central Bank (After the Training Session)

The central bank provides stability to the economy. The bank has one mandate: to stabilize inflation. Inflation is the change of the average price of goods and services in the economy. It is the difference in the price level between the current period and the previous period. The central bank attempts to stabilize inflation at a target of 0 per period. That is, the central bank attempts to make the difference between the current price level and the previous price level 0. If inflation moves higher or lower than 0, due to randomness in the economy, the central bank will act to return inflation back to the target. The bank uses its interest rate to achieve its objective of stabilizing inflation at 0. This implies that the central bank is not concerned with achieving a price-level target but instead attempts to maintain or return inflation to its target. There will be a line on the screen showing you the inflation target.

The Role of the Central Bank Does Not Change. The role of the central bank does not change with this reminder. The simulation will restart, showing the first eight periods exactly as at the start of the session. The simulation is independent of simulation you just completed. Please raise your hand if you have any questions. To continue, the password is the word "continue."

A Reminder of the Role of the Central Bank (Before the Second IT Session)

The central bank provides stability to the economy. The bank has one mandate: to stabilize inflation. Inflation is the change of the average price of goods and services in the economy. It is the difference in the price level between the current period and the previous period. The central bank attempts to stabilize inflation at a target of 0 per period. That is, the central bank attempts to make the difference between the current price level and the previous price level 0. If inflation moves higher or lower than 0, due to randomness in the economy, the central bank will act to return inflation back to the target. The bank uses its interest rate to achieve its objective of stabilizing inflation at 0. This implies that the central bank is not concerned with achieving a price-level target but instead attempts to maintain or return inflation to its target. There will be a line on the screen showing you the inflation target.

The Role of the Central Bank Now Changes (Transition to PLT). The central bank provides stability to the economy. The bank has one mandate: to stabilize the price level. The price level is an average price of goods and services in the economy. The central bank attempts to stabilize the price level at 5 every period. If the price level moves higher or lower than 5, due to randomness in the economy, the central bank will act to return the price level back to the target of 5. The bank uses its interest rate to achieve its objective of stabilizing the price level around the target of 5. This implies that the central bank is not concerned with achieving a constant inflation target but instead attempts to maintain or return the price level to its target. There will be a line on the screen showing you the current price-level target. The simulation will restart, showing the first eight periods exactly as at the start of the session. The simulation is independent of simulation you just completed. You continue to predict inflation each period. The password is "continue."

Figure A.1. A Representative Screen Shown to Experiment Subjects

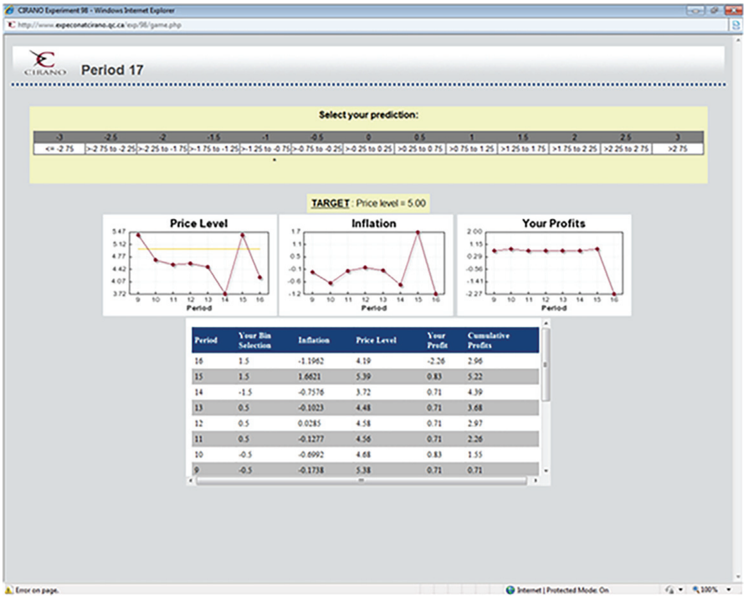

References

Adam, K. 2007. "Experimental Evidence on the Persistence of Output and Inflation." *Economic Journal* 117 (520): 603–36.

Amano, R., O. Kryvtsov, and L. Petersen. 2014. "Recent Developments in Experimental Macroeconomics." *Bank of Canada Review* (Autumn): 1–11.

Anufriev, M., and C. Hommes. 2012. "Evolutionary Selection of Individual Expectations and Aggregate Outcomes in Asset Pricing Experiments." *American Economic Journal: Microeconomics* 4 (4): 35–64.

Arifovic, J., and L. Petersen. 2017. "Stabilizing Expectations at the Zero Lower Bound: Experimental Evidence." *Journal of Economic Dynamics and Control* 82 (September): 21–43.

Assenza, T., P. Heemeijer, C. Hommes, and D. Massaro. 2021. "Managing Self-Organization of Expectations through Monetary Policy: A Macro Experiment." *Journal of Monetary Economics* 117 (January): 170–86.

Bao, T., C. Hommes, and J. Pei. 2021. "Expectation Formation in Finance and Macroeconomics: A Review of New Experimental Evidence." *Journal of Behavioral and Experimental Finance* 32 (December): Article 100591.

Bao, T., C. Hommes, J. Sonnemans, and J. Tuinstra. 2012. "Individual Expectations, Limited Rationality, and Aggregate Outcomes." *Journal of Economic Dynamics and Control* 36 (8): 1101–20.

Berg, C., and L. Jonung. 1999. "Pioneering Price Level Targeting: The Swedish Experience 1931-1937." *Journal of Monetary Economics* 43 (3): 525–51.

Bernanke, B. 2017. "Monetary Policy in a New Era." Prepared for the Rethinking Macroeconomic Policy conference, Peterson Institute, Washington, DC, October 12–13.

Bullard, J. 2018. "Allan Meltzer and the Search for a Nominal Anchor." *Review* (Federal Reserve Bank of St. Louis) 100 (2): 117–26.

Camera, G. 2017. "A Perspective on Electronic Alternatives to Traditional Currencies." *Sveriges Riksbank Economic Review* (1): 126–48.

Cateau, G., and M. Shukayev. 2022. "Limited Commitment, Endogenous Credibility and the Challenges of Price-Level Targeting." *Canadian Journal of Economics* 55 (4): 1834–61.

Coibion, O., and Y. Gorodnichenko. 2011. "Monetary Policy, Trend Inflation, and the Great Moderation: An Alternative Interpretation." *American Economic Review* 101 (1): 341–70.

Coibion, O., Y. Gorodnichenko, E. Knotek II, and R. Schoenle. 2023. "Average Inflation Targeting and Household Expectations." *Journal of Political Economy Macroeconomics* 1 (2): 403–46.

Cornand, C., and F. Heinemann. 2014. "Experiments on Monetary Policy and Central Banking." In *Experiments in Macroeconomics,* Vol. 17, ed. J. Duffy, 167–227. Emerald Group Publishing Limited.

———. 2019. “Experiments in Macroeconomics: Methods and Applications.” In *Handbook of Research Methods and Applications in Experimental Economics*, ed. A. Schram and A. Ule, 269–94. Edward Elgar Publishing.

Cornand, C., and P. Hubert. 2020. “On the External Validity of Experimental Inflation Forecasts: A Comparison with Five Categories of Field Expectations.” *Journal of Economic Dynamics and Control* 110 (January): Article 103746.

———. 2022. “Information Frictions across Various Types of Inflation Expectations.” *European Economic Review* 146 (July): Article 104175.

Cornand, C., and C. M’baye. 2018a. “Band or Point Inflation Targeting? An Experimental Approach.” *Journal of Economic Interaction and Coordination* 13 (2): 283–309.

———. 2018b. “Does Inflation Targeting Matter? An Experimental Investigation.” *Macroeconomic Dynamics* 22 (2): 362–401.

Duffy, J. 2017. “Macroeconomics: A Survey of Laboratory Research.” In *Handbook of Experimental Economics*, Vol. 2, ed. J. Kagel and A. Roth, 1–90. Princeton, NJ: Princeton University Press.

Eggertsson, G., and M. Woodford. 2003. “The Zero Bound on Interest Rates and Optimal Monetary Policy.” *Brookings Papers on Economic Activity* 1: 139–211.

Engle-Warnick, J., and N. Turdaliev. 2010. “An Experimental Test of Taylor-type Rules with Inexperienced Central Bankers.” *Experimental Economics* 13 (2): 146–66.

Evans, C. 2010. “Monetary Policy in a Low-Inflation Environment: Developing a State-Contingent Price-Level Target.” Remarks delivered before the Federal Reserve Bank of Boston’s 55th Economic Conference, October 16.

Galí, J., and M. Gertler. 1999. “The Science of Monetary Policy: A New Keynesian Perspective.” *Journal of Economic Literature* 37 (4): 1661–707.

Greiner, B. 2004. “The Online Recruitment System ORSEE 2.0 — A Guide for the Organization of Experiments in Economics.” Working Paper Series in Economics No. 10, University of Cologne.

Hebden, J., E. Herbst, J. Tang, G. Topa, and F. Winkler. 2020. "How Robust Are Makeup Strategies to Key Alternative Assumptions?" Finance and Economics Discussion Series No. 2020-069, Board of Governors of the Federal Reserve System.

Hommes, C. 2011. "The Heterogeneous Expectations Hypothesis: Some Evidence from the Lab." *Journal of Economic Dynamics and Control* 35 (1): 1–24.

Hommes, C., and T. Makarewicz. 2021. "Price Level versus Inflation Targeting under Heterogeneous Expectations: A Laboratory Experiment." *Journal of Economic Behavior and Organization* 182 (February): 39–82.

Hommes, C., D. Massaro, and M. Weber. 2019. "Monetary Policy under Behavioral Expectations: Theory and Experiment." *European Economic Review* 118 (September): 193–212.

Honkapohja, S., and K. Mitra. 2020. "Price Level Targeting with Evolving Credibility." *Journal of Monetary Economics* 116 (December): 88–103.

Huber J., M. Kirchler, and T. Stockl. 2010. "The Hot Hand Belief and the Gambler's Fallacy in Investment Decisions under Risk." *Theory and Decision* 68 (4): 445–62.

Kostyshyna, O., L. Petersen, and J. Yang. 2022. "A Horse Race of Monetary Policy Regimes: An Experimental Investigation." Staff Working Paper No. 2022-33, Bank of Canada.

Kryvtsov, O., and L. Petersen. 2019. "Expectations and Monetary Policy: Experimental Evidence." Manuscript, Simon Fraser University, Vancouver, BC, Canada.

———. 2021. "Central Bank Communication that Works: Lessons from Lab Experiments." *Journal of Monetary Economics* 117 (January): 760–80.

Lambsdorff, J., M. Schubert, and M. Giamattei. 2013. "On the Role of Heuristics—Experimental Evidence on Inflation Dynamics." *Journal of Economic Dynamics and Control* 37 (6): 1213–29.

Luhan, W., and J. Scharler. 2014. "Inflation Illusion and the Taylor Principle: An Experimental Study." *Journal of Economic Dynamics and Control* 45 (August): 94–110.

Meh, C., J. Ríos-Rull, and Y. Terajima. 2010. "Aggregate and Welfare Effects of Redistribution of Wealth under Inflation and Price-Level Targeting." *Journal of Monetary Economics* 57 (6): 637–52.

Mirdamadi, M., and L. Petersen. 2019. "Macroeconomic Literacy and Expectations." Manuscript, Simon Fraser University, Vancouver, BC, Canada.

Newey, W., and K. West. 1987. "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55 (3): 703–8.

Petersen, L. 2014. "Forecast Error Information and Heterogeneous Expectations in Learning-to-Forecast Experiments." In *Experiments in Macroeconomics,* Vol. 17, ed. J. Duffy, 109–37. Emerald Group Publishing Limited.

Petersen, L., and R. Rholes. 2022. "Macroeconomic Expectations, Central Bank Communication, and Background Uncertainty: A COVID-19 Laboratory Experiment." *Journal of Economic Dynamics and Control* 143 (October): Article 104460.

Pfajfar, D., and B. Žakelj. 2014. "Experimental Evidence on Inflation Expectation Formation." *Journal of Economic Dynamics and Control* 44 (July): 147–68.

———. 2018. "Inflation Expectations and Monetary Policy Design: Evidence from the Laboratory." *Macroeconomic Dynamics* 22 (4): 1035–75.

Roos, M., and W. Luhan. 2008. "Are Expectations Formed by the Anchoring-and-Adjustment Heuristic? An Experimental Investigation." Ruhr Economic Paper No. 54.

Salle, I. 2023. "What to Target? Insights from Theory and Lab Experiments." *Journal of Economic Behavior and Organization* 212 (August): 514–33.

Svensson, L. 1999. "Price-Level Targeting versus Inflation Targeting: A Free Lunch?" *Journal of Money, Credit and Banking* 31: 277–95.

———. 2003. "Escaping From a Liquidity Trap and Deflation: The Foolproof Way and Others." *Journal of Economic Perspectives* 17 (4): 145–66.

Tversky, A., and D. Kahnemann. 1974. "Judgement under Uncertainty: Heuristics and Biases." *Science* 185 (4157): 1124–32.

Vestin, D. 2006. "Price-Level versus Inflation Targeting." *Journal of Monetary Economics* 53 (7): 1361–76.

Williams, J. 2018. “The Options: Keep It, Tweak It, or Replace It.” Panel discussion at conference entitled “Should the Fed Stick with the 2 Percent Inflation Target or Rethink It?” held at Brookings Institute, Washington, DC, January 8.

Woodford, M. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton, NJ: Princeton University Press.