Imperfect Macroeconomic Expectations: Evidence and Theory*

George-Marios Angeletos† **Zhen Huo**‡ **Karthik A. Sastry**§

April 2, 2020

[preliminary and incomplete]

Abstract

We complement existing evidence about the predictability of average and individual forecast errors with a new fact: in response to aggregate shocks, expectations under-react initially but over-shoot later on. We next inspect these facts under the lens of a parsimonious theoretical framework that distills the essence of diverse existing theories about expectation formation. We conclude that the evidence favors a theory that combines dispersed information with over-extrapolation. Theories that emphasize cognitive discounting, level-k thinking and over-confidence find little support. We finally illustrate the general-equilibrium implications of the documented facts within the New Keynesian model.

*This paper is under preparation for the 2020 NBER Macroeconomics Annual. We are thankful to our discussants, Ricardo Reis and Jessica Watcher, and the organizers, Martin Eichenbaum and Eric Hurst. We also thank Fabrice Collard for help with some of the material in Section 4. Angeletos acknowledges the financial support of the National Science Foundation (Award #1757198).

†MIT and NBER; angelet@mit.edu

‡Yale University; zhen.huo@yale.edu

§MIT; ksastry@mit.edu

1 Introduction

The rational expectations hypothesis is one of the most influential and widely applied ideas in macroeconomics. It is often combined with a strong, complementary hypothesis that the information on which these rational expectations are based is common to all people. But an explosion of recent theoretical and empirical work has questioned both premises. This has pushed the discipline back toward reckoning with a "wilderness" of alternative models for expectations formation and equilibrium.

What does survey evidence on expectations tell us within the space of the alternative hypotheses? Are "imperfect expectations" due to informational frictions, systematic biases in beliefs, or both? And which kind of evidence is most useful for gauging the quantitative bite of belief frictions?

In the hopes of answering these questions, and helping identify "where we are in the wilderness," this article attempts to draw some of the recent theoretical and empirical literature on macroeconomic expectations under a common umbrella. We develop a parsimonious framework that allows for both informational frictions and mis-specified beliefs (or bounded rationality). We use this framework to organize existing survey evidence on expectations, guide the discovery of new evidence, and ultimately pin down the "right" model of expectations for the business-cycle context.

Previous empirical studies have also tried to disentangle mechanisms from surveys of expectations.[1] But we will argue that our own empirical strategy, with its focus on dynamic impulse responses to shocks that account for the bulk of the business-cycle variation in unemployment and inflation, provides sharper guidance. This positive contribution dovetails with our offering guidance about how to interpret, and *use*, the evidence in a general equilibrium context, where expectations and outcomes feed to one another.

We find that the right model combines dispersed information with over-extrapolation. Following any innovation, the informational friction is the dominant force initially, helping explain not only why average forecasts under-react on impact but also why one's forecast error is predictable by the past revisions of *others*. But as time passes and agents learn, over-extrapolation eventually takes over. This helps explain why average forecast errors reverse sign after a few quarters—a fact that favors over-extrapolation over under-extrapolation and the latter's close cousins, cognitive discounting and level-K thinking.

A flexible framework. The theoretical framework we use in this paper is an extension of that developed in Angeletos and Huo (2019). Our extension is both highly tractable and highly parsimonious: just three parameters govern beliefs away from the full-information, rational-expectations benchmark. And yet, it is flexible enough to capture the essence of a large literature on imperfect macroeconomic expectations.

One strand of this literature emphasizes dispersed private information (Lucas, 1972; Lorenzoni, 2009), rational inattention (Sims, 2003; Mackowiak and Wiederholt, 2009), and sticky information (Mankiw and Reis, 2002). Another strand emphasizes the importance of strategic complementarity and higher-order uncertainty in such environments (Morris and Shin, 2002; Woodford, 2003; Nimark, 2008; Angeletos and Lian, 2016, 2018). Moving beyond the rational model, some authors emphasize biases to over-extrapolate the past (Gennaioli, Ma, and Shleifer, 2015a; Fuster, Laibson, and Mendel, 2010; Guo and Wachter, 2019), or

[1] e.g., Coibion and Gorodnichenko (2012, 2015); Kohlhas and Broer (2019); Bordalo et al. (2018); Fuhrer (2018)

"cognitively discount" the future (Gabaix, 2019). An emerging macroeconomic literature on level-K thinking also boils down to discounting of the future (Garcıa-Schmidt and Woodford, 2019; Farhi and Werning, 2019). A final strand emphasizes under- or overconfidence in various information sources (Kohlhas and Broer, 2019), or prioritization of those that seem "representative" (Bordalo, Gennaioli, and Shleifer, 2017).

Because of the desire for parsimony, our framework does not give full justice to the entire richness of this diverse set of theories. Instead, it distills their most essential ingredients, those that drive their business cycle implications and their empirical footprint on expectations.

Another quality of our framework is the accommodation of the equilibrium fixed point between expectations and behavior. This is paramount for quantifying the causal effect of belief frictions on macroeconomic outcomes. But it is not strictly needed for organizing the empirical evidence: for this purpose, it is fine to treat the macroeconomic outcomes as exogenous processes and focus on the forecasting problem of the individuals. This motivates the step-by-step approach described next.

Evidence. In Sections 3 and 4, we revisit two existing facts about the predictability of forecast errors and provide a new, third fact about their dynamic response to shocks.

F1. For both unemployment and inflation, aggregate forecast errors are positively related to lagged aggregate forecast revisions, as in Coibion and Gorodnichenko (2015). This pattern suggests that aggregate forecasts *under-react* to aggregate news.

F2. The corresponding individual-level relation, previously explored in Bordalo et al. (2018) and Kohlhas and Broer (2019), is more complicated. In the case of inflation, individual forecasts appear to *over-react* to own revisions, in sharp contrast to the corresponding aggregate fact. And in the case of unemployment, they *under-react*, as in the aggregate evidence but with less ferocity.

F3. Consider two semi-structural shocks, one that accounts for most of the business-cycle variation in unemployment and other macroeconomic quantities, and another that accounts for most of the business-cycle variation in inflation.[2] Construct the Impulse Response Functions (IRFs) of the average *forecasts* of unemployment and inflation to the corresponding shocks. In both cases, average forecasts are initially under-react before over-shooting later on, or predicting larger and longer-lasting effects of the shock than those that occur.

Interpretation. In Section 5, we offer a structural interpretation of the aforementioned facts under the lens of a simplified version of the framework, which abstracts from the fixed-point relation between expectations and outcomes. This makes the exercise directly comparable to the related empirical literature and lets us extract the relevant information from the data in the most transparent manner.

We first explain why Facts 1 and 2, by themselves, cannot discern the distortions in beliefs. For instance, Coibion and Gorodnichenko's (2015) interpretation of Fact 1 as a measure of the informational friction is invalid when there is a departure from rational expectations. Similarly, Kohlhas and Broer (2019) and Bordalo et al. (2018) use Fact 2 to argue that the requisite departure from rational expectations is

[2] These shocks are described at the end of Section 2 and are obtained from Angeletos, Collard, and Dellas (2019).

overconfidence (over-estimating the precision of one's information), or a close cousin of it ("representativeness"). But the same fact could also be consistent with over-extrapolation (over-estimating the persistence of the underlying process). Furthermore, the evidence itself is not clear-cut: the unemployment forecasts exhibit the opposite individual-level pattern than the inflation forecasts.

We next explain how Fact 3 *alone* helps nail down the "right" combination of frictions: to match this fact, it is *necessary* to combine slow learning with over-extrapolation, regardless of the degree of over- or under-confidence. We provide additional, more direct evidence of over-extrapolation by showing that the subjective persistence, as revealed by the term structure of expectations, is larger than the objective persistence, as measured by the impulse response of the outcome. We finally show how the three facts together help identify the precise quantitative combination of the underlying belief parameters.

GE and application to New Keynesian model. In Section 6, we incorporate a GE feedback between expectations and outcomes. We explain how the fixed point works and highlight how it depends not only the considered frictions in beliefs but also on parameters that determine the relative strength of PE and GE effects.[3] On the one hand, this helps us underscore how the bite of the belief distortions hinges of GE multipliers and policy. On the other hand, it allows us to illustrate the tight connection between under-extrapolation, cognitive discounting, and levek-K thinking—and to extend the aforementioned lesson about the "right" model of beliefs to a GE context.

In Section 7, we finally illustrate the "bite" of imperfect expectations within the three-equation New Keynesian model. We do so by mapping this bite to the survey evidence documented in the first part of our paper. We explain how the dynamic properties of the *average* forecasts we emphasize in this paper can serve as "sufficient statistics" for the counterfactuals of interest, leaving the properties of the *individual* forecasts emphasized in Bordalo et al. (2018) to an afterthought.

Discussion. The evidence we marshal suggests that a combination of noisy information (or imprecise perception), slow learning, and over-extrapolation best describes the "business cycle behavior" of survey expectations for unemployment and inflation. This is a positive result, which generates new, and quantifiable, predictions for the actual macroeconomic dynamics. But it also reveals two slightly more "negative" lessons about what models are *not* ideal for fitting the evidence and, symmetrically, which evidence is *not* ideal for informing macroeconomic theory.

The first lesson concerns the behavioral interpretation of persistence. What we can rule out as good models of macro expectations "on average" are those that rely heavily on under-extrapolation of the present to the future. The theories of cognitive discounting and level-K thinking have this property. We argue they are not ideally suited to standard business cycle analysis for a slightly subtle reason that emerges only after we try to calibrate to many different moments at once—while macro expectations are indeed sluggish and myopic, they do seem to over-extrapolate the effects of a *given* shock over time. This conclusion need not invalidate such models in inherently *non-stationary* regimes, including the time at the ZLB during the

[3] Echoing lessons from Angeletos and Lian (2018), Angeletos and Huo (2019), and Farhi and Werning (2019), such parameters include the subjective discount factor and the slope of the Keynesian cross—objects that drop out in the full-information, rational-expectations benchmark but are crucial away from it, for they determine the importance of higher-order beliefs.

Great Recession and the abrupt rate changes of the Volcker disinflation.[4] But it does suggest that such theories do not capture the behavior of expectations during more usual fluctuations.

The second lesson concerns the appropriate interpretation of individual-level rejections of rational expectations (like those presented by Bordalo et al., 2018; Kohlhas and Broer, 2019; Fuhrer, 2018). Our overall finding, from combining individual and aggregate moments, is that the combination of noise and over- and under-extrapolation is a more compelling theory than "overconfidence" in one's own projections. Once we allow for the kind of over-extrapolation needed to account for the dynamic response of the average forecasts, the evidence on individual forecast error predictability seem to require *underconfidence*. How this finding should relate to the theoretical and experimental literatures on subjective confidence in beliefs is beyond the scope of this article. But we also argue that such a decomposition of individual deviations into "extrapolation" versus "confidence" effects, while appropriate for understanding the complete landscape of expectations formation, could be inconsequential for the main counterfactual of interest.

Other related literature. The macroeconomics literature on informational frictions and non-rational expectations is large and growing. A few examples among the many issues we abstract from are: the endogeneity of market signals (e.g. Amador and Weill, 2010; Chahrour and Gaballo, 2018; Hassan and Mertens, 2014); information choice (e.g., Mackowiak and Wiederholt, 2009, 2015; Sims, 2010; Veldkamp, 2006, 2011); the macro-finance implications of heterogeneous or mis-specified beliefs (e.g., Caballero and Simsek, 2017; Geanakoplos, 2010; Simsek, 2013); noise- or sentiment-driven fluctuations (e.g., Angeletos and La'O, 2013; Benhabib, Wang, and Wen, 2015; Lorenzoni, 2010); non-rational belief contagion (e.g., Burnside, Eichenbaum, and Rebelo, 2016; Carroll, 2001); robustness and ambiguity (e.g., Ilut and Schneider, 2014; Bhandari, Borovička, and Ho, 2019); influence of experiences and socioeconomic status on individual expectations formation (e.g., Das, Kuhnen, and Nagel, 2020; Malmendier and Nagel, 2016; D'Acunto et al., 2019); and adaptive expectations (e.g., Eusepi and Preston, 2011; Evans and Honkapohja, 2012; Sargent, 2001). Instead, we opt to focus on what, at least in our view, are the common threads of the particular strands of the theoretical literature we mentioned earlier and to connect them with the emerging survey evidence on expectations.

Related to our evidence about dynamic responses (Fact 3) above is an earlier literature documenting the pattern of slow initial reaction and subsequent over-reaction in individual stock prices (De Bondt and Thaler, 1985; Cutler, Poterba, and Summers, 1991; Lakonishok, Shleifer, and Vishny, 1994). Theoretical work such as Barberis, Shleifer, and Vishny (1998), Daniel, Hirshleifer, and Subrahmanyam (1998), and Hong and Stein (1999) provide parsimonious interpretations which combine tentative initial reactions (either at the individual or group level) with medium-run over-reaction to news. More recently, Greenwood and Shleifer (2014) and Gennaioli, Ma, and Shleifer (2015b) demonstrate patters in survey expectations of stock returns and firm earnings that are also suggestive of over-extrapolation. However, all these works do not share our focus on business cycles and, most importantly, do not provide the kind of evidence we offer about the dynamics responses (IRFs) of forecast errors to shocks and the term structure of forecasts.

[4] For studies of the former topic that rely on variations of "under-extrapolation," see Farhi and Werning (2019), Gabaix (2019), Iovino and Sergeyev (2017), Garcıa-Schmidt and Woodford (2019), and Angeletos and Sastry (2020).

Roadmap. We review data and measurement in Section 2. We present our "three facts" in Sections 3 and 4 and interpret them in a simplified, non-GE version of our model in Section 5. We show in Section 6 how to integrate the obtained lessons tractably into a GE context and how to connect to various strands of the literature. We illustrate the "bite" of the documented frictions within the New Keynesian model in Section 7. We conclude with a discussion of our findings and the implications for future work in Section 8.

2 Data and Measurement

We will focus on two macroeconomic outcomes: unemployment and inflation. We now review the exact data sources we use for forecasts and realized outcomes of these variables.

Forecasts from the Survey of Professional Forecasters. Our main dataset for forecasts is the Survey of Professional Forecasters, a panel survey of about 40 experts from industry, government, and academia, currently administered by the Federal Reserve Bank of Philadelphia. Every quarter, each survey respondent is asked for point-estimate projections of several macro aggregates. Our main sample runs from Q4 1968 to Q4 2017. We will use forecasts of the civilian unemployment rate (averaged over the quarter) and the growth rate in GDP deflator.

At various points in the analysis, we will require "consensus" and "individual-level" forecasts of each macro variable. For the "consensus," we always use the *median* forecast of the object of interest (e.g., unemployment or inflation at a given horizon). Using the median instead of the mean alleviates concerns about outliers and/or data-entry errors, which could be quite influential in the 40-forecaster cross section, from driving the results.

For the individual-level results, we always trim outliers in forecast errors and revisions that are plus or minus 4 times the inter-quartile range from the median, where both reference values are calculated over the entire sample. [5]

Other survey sources. We also provide corroborating evidence from two additional survey datasets. The first is the Blue Chip Economic Indicators Survey, a privately-operated professional forecast with a similar scale and scope to the SPF. We use Blue Chip data from 1980 to 2017 and focus on the reported "consensus forecast" for unemployment and GDP deflator. This dataset is available at the monthly frequency, so we use end-of-quarter forecasts (i.e., those made in March, June, September, and December) for comparability with the SPF.

The second source is the University of Michigan Survey of Consumers, which is (for our purposes) a repeated cross-section of about 500 members of the "general public" contacted by phone.[6] Like with the Blue Chip survey, we focus on end-of-quarter waves. We take the Michigan survey inflation forecast

[5] For context, in a Gaussian distribution, the probability of an observation so far in the tails is about 6.8×10^{-8}. Nonetheless, in the sample of three-quarter ahead inflation forecast errors, there are 57 such observations out of 7,438 forecaster-quarter observations, or 112,408 times the expectation were the data Gaussian. To give a sense of magnitude (and "plausibility" of making such errors), all of the outliers involve forecast errors greater than 5.37 percentage points in the estimate of annual inflation.

[6] There is a panel component of the Michigan survey, in which some respondents are re-contacted after 6 months, but this does not help us calculate forecast revisions since the forecasting horizon has changed so much.

as the median response to the question about price increases.[7] We code also a forecast for the growth rate of unemployment based on a question about whether unemployment will increase or decrease over the coming twelve months.[8] For this measure we take the cross-sectional mean, which corresponds to a "consensus forecast" about the sign of the growth rate of unemployment.

Macro data (and vintages thereof). Our unemployment measure u_t will be the average BLS unemployment rate in a given quarter t. Our inflation measure $\pi_{t,k}$ will be the percentage increase in GDP or GNP deflator between period t and period $t-k-1$.[9] The timing assumption matches the fact that forecasters have access to (first-revision) data for inflation and the price level at $t-k-1$ when forecasting k-periods ahead at $t-k$. We will let outcome x_{t+k} refer to u_{t+k}, unemployment k periods ahead, or $\pi_{t+k,k}$, the inflation from the reference period $t-1$ to the future period $t+k$.

In our replication of Coibion and Gorodnichenko (2015) and Bordalo et al. (2018) in Sections 3.1-3.2, we will use *first-vintage* macro data to measure actual outcomes, because this is what these papers did in the first place.[10] However, such measurement is not necessarily the right one from the perspective of theory. If agents are try to forecast the *actual* levels of unemployment and inflation, the "econometrician" should use the *final-release* data. We will thus verify the robustness of the relevant exercises to the use of final-release data.

Finally, in our study of IRFs in Section 4, we will use *final-release* data both for the aforementioned reason and for consistency with the main macro time-series literature. But once again, we will consider the "opposite" measurement (in this case, first-vintage data in place of final-release data) for robustness.

Shocks. Finally, we use two semi-structurally identified shocks from Angeletos, Collard, and Dellas (2019). The empirical strategy taken in that paper builds on the max-share approach (Uhlig, 2003; Barsky and Sims, 2011) but is guided by the following goal: providing a parsimonious representation of the business cycle in terms of one or a dominant business cycle shock.

To this goal, Angeletos, Collard, and Dellas (2019) run a VAR on a set of ten or more key macroeconomic variables that includes the two variables we focus on here, the rate of unemployment and the rate of inflation.[11] They then compile a collection of multiple shocks, each identified by maximizing its contribution to the volatility of a particular variable over a particular frequency band, and they draw lesson from comparing the empirical footprint of all these shocks.

[7] The exact question is the following: "By about what percent do you expect prices to go (up/down) on the average, during the next 12 months?" Respondents can key in a response rounded to the nearest whole number.

[8] The exact question is the following: "How about people out of work during the coming 12 months. Do you think that there will be more unemployment than now, about the same, or less?" There are three responses, as indicated in the question.

[9] The ambiguity between GDP and GNP matches the fact that the Survey of Professional Forecasters changed its main target variable from GNP (and the deflator thereof) to GDP (and the deflator thereof) starting in 1992.

[10] We take all vintage data series from the Philadelphia Fed's website: https://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/

[11] The other eight variables in their baseline VAR are: the per-capita levels of GDP, investment (inclusive of consumer durables), consumption (of non-durables and services), and total hours worked; labor productivity in the non-farm business sector; utilization-adjusted TFP; the labor share; and the federal funds rate. Larger VARs that add stock prices and credit indicators yield similar results.

Consider the shock that is constructed by maximizing its contribution to the variation in unemployment at business-cycle frequencies (6-32 quarters). This shock is found to have the following properties. First, it is nearly identical, in terms of IRFs and variance contributions, to the shocks that target the business-cycle variation in any of the following other variables: hours worked, GDP, investment, consumption, and the output or unemployment gap. Second, it captures the majority of the business-cycle variation in all these variables, and strong positive co-movement among them. Third, it has a negligible footprint on TFP at all horizons. And finally, it has a small to modest footprint on inflation.

These facts together provide support for parsimonious theories that attribute the bulk of the business cycle to a single, non-inflationary or mildly-inflationary, demand shock. They also motivate us to consider, in our empirical exercises, a single-shock representation of the joint dynamics of actual unemployment and the forecasts thereof, where the underlying driving force is the aforementioned "unemployment shock" from Angeletos, Collard, and Dellas (2019)

The second shock we borrow from Angeletos, Collard, and Dellas (2019) is identified by maximizing its contribution to the business cycle variation in inflation. This shock accounts for over 80% of the business-cycle variation in inflation, but is is found to have a small footprint on all real quantities, including TFP. It can thus be interpreted as some kind of non-technology supply shock, or a markup shock, which manifests primarily in inflation. And it motivates us to consider, in our empirical exercises, a single-shock representation of the joint dynamics of actual inflation and the forecasts thereof, where the underlying driving force is this "inflation shock."

Whether these shocks, or any other SVAR-based shocks, are "truly" structural is beyond the scope of the present paper. For our purposes, the appeal of these particular shocks compared to others found in the literature is that they drive a significant component of the business cycle variation in macroeconomic activity and inflation. There is thus a good chance that they also drive a significant component of the corresponding variation in people's expectations.

3 Two (Unconditional) Facts About Forecasts

This section presents two facts about macro survey forecasts and forecast errors, which relate to *unconditional moments* and are known from the literature. They concern the correlation between forecast errors and previous-period forecast revisions at the average level (Fact 1) and the individual level (Fact 2). The final fact (Fact 3), which is new, concerns the IRFs of average forecasts and average forecast errors to macroeconomic shocks, and in particular the "reversal of sign" of dynamic response of the errors at medium horizons. We will first present all the facts "as is," and then in Section 5 suggest a unifying theoretical explanation.[12]

[12]The empirical literature on forecasts is large and growing. The papers we build more heavily on are Coibion and Gorodnichenko (2012, 2015), Bordalo et al. (2018), Kohlhas and Broer (2019), and Fuhrer (2018); the connections are explained below. An earlier important contribution is Mankiw, Reis, and Wolfers (2004); but this paper focuses on the cross-section dispersion of forecast errors, a moment that is of little use for our purposes for reasons explained in due course.

Table 1: Predicting Aggregate Forecast Errors

$$\text{Error}_t = K_{CG} \cdot \text{Revision}_t + \alpha + u_t$$

	(1)	(2)	(3)	(4)
	Unemployment		Inflation	
sample	1968-2017	1984-2017	1968-2017	1984-2017
Revision_t (K_{CG})	0.741	0.809	1.528	0.292
	(0.232)	(0.305)	(0.418)	(0.191)
R^2	0.111	0.159	0.278	0.016
Observations	191	136	190	135

Notes: The dataset is the Survey of Professional Forecasters and the observation is a quarter between Q4-1968 and Q4-2017. All regressions include a constant. The forecast horizon is 3 quarters. Standard errors are HAC-robust, with a Bartlett ("hat") kernel and lag length equal to 4 quarters. The data used for outcomes are first-release ("vintage").

3.1 Fact 1: Aggregate errors are predictable by aggregate past revisions

Coibion and Gorodnichenko (2015), henceforth CG, test for a departure from full-information rational expectations (or FIRE for short) by estimating the predictability of professionals' aggregate ("consensus") forecast errors using information in previous forecast-revisions.

Let $\bar{\mathbb{E}}_t[x_{t+k}]$ denote the average or median expectation of variable x_{t+k} (either unemployment or inflation) measured at time t. Let $\bar{\mathbb{E}}_{t-1}[x_{t+k}]$ be the median forecast at time $t-1$.[13] The associated *forecast error* from time t is $\text{Error}_{t,k} \equiv x_{t+k} - \bar{\mathbb{E}}_t[x_{t+k}]$, suppressing notation for what variable x is being forecast, and the *forecast revision* is $\text{Revision}_{t,k} \equiv \bar{\mathbb{E}}_t[x_{t+k}] - \bar{\mathbb{E}}_{t-1}[x_{t+k}]$.

CG run the following regression that projects aggregate forecast errors onto aggregate forecast revisions:

$$\text{Error}_{t,k} = \alpha + K_{\text{CG}} \cdot \text{Revision}_{t,k} + u_{k,t} \tag{1}$$

where K_{CG}, in shorthand notation that references the authors, is the main object of interest.

The fact: $K_{CG} > 0$ for both inflation and unemployment. Table 1 reports results from estimating (1) at the horizon $k = 3$ for both unemployment and inflation in our data. We report results over the full sample 1968-2017 (columns 1 and 3), and also over a restricted sample after 1984 (columns 2 and 4). We may believe *a priori* that the latter is a more consistent and "stationary" regime for the US macroeconomy (i.e., after the oil crisis and Volcker disinflation).

Like Coibion and Gorodnichenko (2015), we find in all specifications a point estimate of $K_{\text{CG}} > 0$: when professional forecasters, in aggregate, revise upward their estimation of unemployment or inflation, they on average always "undershoot" the eventual truth. For inflation, we find the predictability is considerably lower on the restricted sample. Appendix Tables A.1 and A.2 show the results for all horizons over the full sample and post-1984 sample, respectively.

[13] In the data, we prefer to use the median to limit the influence of outliers and/or data entry errors. But results with the mean are essentially identical. In the theory, means and medians coincide because we let all variables and signals be Normally distributed.

Interpretation. Suppose we approach regression (1) through the perspective of FIRE, that is, with a model featuring a representative, rational agent. Under this perspective, $\text{Error}_{t,k}$ and $\text{Revision}_{t,k}$ are empirical proxies of that single agent's realized forecast error and past forecast revisions, respectively. In theory, the forecast error has to be unpredictable by past information, and hence also by past revisions. It follows that, unless the measurement error in these proxies is correlated, K_{CG} ought to be zero.

We just saw that $K_{\text{CG}} > 0$ in the data. This finding rejects FIRE and gives some guidance on what kind of departures from that benchmark one should contemplate. In particular, the following *literal* interpretation of the sign of K_{CG} is valid: whenever forecasts adjust upward today, forecast errors tomorrow tend to be positive, or forecasts should have adjusted upward even more to accurately track the actual outcome. In other words, aggregate forecasts too "sluggishly" relative to FIRE.

Coibion and Gorodnichenko (2015) interpret this finding as evidence of noisy, dispersed information. Indeed, models such as those articulated in Mankiw and Reis (2011), Woodford (2003), and Nimark (2008), naturally give rise the aforementioned empirical pattern. But it is important to keep in mind the following two qualifications.

First, what is key for the ability of such models to generate $K_{\text{CG}} > 0$ is the assumption that information is not only *noisy* but also *dispersed*, or heterogeneous: when information is noisy but commonly shared, rational expectations imposes $K_{\text{CG}} = 0$ regardless of what that information is. By the same token, rational inattention (Sims, 2003) helps accommodate $K_{\text{CG}} > 0$ only insofar rational inattention is a micro-foundation of dispersed private information.

Second, although dispersed, noisy information is sufficient for $K_{\text{CG}} > 0$, it is not necessary: the same fact could also mean a departure from rational expectations. For instance, adaptive expectations, underextrapolation, cognitive discounting, and level-K thinking can also generate $K_{\text{CG}} > 0$. We will clarify the mapping from K_{CG} to these theories in Sections 5 and 6, and we will show how the combination of $K_{\text{CG}} > 0$ with the additional two facts reported in the rest of this section help select the "right" explanation among the multiple candidates allowed.

Notwithstanding these points, what *is* clear at this point is that $K_{\text{CG}} > 0$ is, even by itself, is a rejection of "perfect expectations."

A business-cycle-IV version of the CG regression . One immediate limitation of estimating (1) is that it makes no distinction between sources of variation in forecast revisions—that is, the business cycle variation that matters in models is treated equally with measurement corrections, seasonal fluctuations, and other "less interesting" sources of variation. As one simple alternative, we estimate a version of (1) with the forecast revision instrumented by the current value and six lags of the variable-specific "business-cycle shock" from Angeletos, Collard, and Dellas (2019). This allows us to zoom in on the forecast revision variation that is germane to main business-cycle fluctuations in unemployment and GDP deflator, respectively.

Table 2 shows that, for unemployment, this estimate tends to agree with the OLS or suggest weakly that predictability has increased over the modern sample; while for inflation, it agrees with OLS on the full sample and provides a much larger estimate on the modern sub-sample. This last point suggests that much of the difference between sub-samples in Table 1 relates to the *composition of shocks* rather than

Table 2: Predicting Aggregate Forecast Errors: An IV Approach

$$\text{Error}_t = K_{CG} \cdot \text{Revision}_t + \alpha + u_t$$
$$\text{Revision}_t = \rho \cdot \text{Shock}_t + \phi + u_t$$

	(1)	(2)	(3)	(4)
	Unemployment		Inflation	
sample	1968-2017	1984-2017	1968-2017	1984-2017
Revision_t (K_{CG})	0.585	0.983	1.460	1.038
	(0.393)	(0.264)	(0.521)	(0.467)
Instruments	UnempShock_{t-h}: 0≤h≤6		InflShock_{t-h}: 0≤h≤6	
First-stage F	7.527	4.736	3.517	5.047
OLS Estimate	0.741	0.809	1.528	0.292
Observations	189	130	188	130

Notes: The dataset is the Survey of Professional Forecasters and the observation is a quarter between Q4-1968 and Q4-2017. All regressions include a constant. The forecast horizon is 3 quarters. Standard errors are HAC-robust, with a Bartlett ("hat") kernel and lag length equal to 4 quarters. The data used for outcomes are first-release ("vintage").

the conditional response of inflation and forecasts thereof to specific shocks. This will be one of multiple motivating reasons for us to switch the focus to *conditional shock responses* later in the empirical analysis.

Robustness. With an eye looking ahead to calibrating macroeconomic models, it is not clear whether we want to focus on the first-release data releases for outcomes rather than the final-release data on outcomes, which may be purged from measurement errors and otherwise be a more correct gauge of economic activity. The latter is preferable if forecasters are trying to forecast the "truth," and seems more appropriate vis-a-vis the theories we consider in this paper. For this reason, Table A.3 recreates the CG evidence using the final release data for the outcomes. The results are extremely similar.

Finally, Table A.4 in the Appendix recreates the CG evidence using Blue Chip Economic Indicators data (1980-2017). This uncovers largely the same patterns, especially compared to the later sub-sample.

3.2 Fact 2: Individual errors are predictable by own past revisions

Models with noisy and dispersed information, in their simplest form, retain individual-level rationality. As such, these models allow for predictability of forecast errors at the *aggregate* level but rule out such predictability at the *individual* level. That is, a given agent's forecast error should be predictable by *their own* past information.

Recent papers by Bordalo et al. (2018), Fuhrer (2018), and Kohlhas and Broer (2019) have turned to individual-level, panel data on forecasts to test such a prediction. All three estimate an analogue to (1) at the individual level. Bordalo et al. (2018) contain a relatively more extensive exploration across forecasts of different objects (inflation, unemployment, output, and various interest rates). For our purposes, we focus on forecasts of inflation and unemployment.

Let $\text{Error}_{i,t,k} \equiv x_{t+k} - \mathbb{E}_{i,t}[x_{t+k}]$ and $\text{Revision}_{i,t,k} \equiv \mathbb{E}_{i,t}[x_{t+k}] - \mathbb{E}_{i,t-1}[x_{t+k}]$ denote forecast errors and

Table 3: Predicting Individual Forecast Errors

$$\text{Error}_{i,t} = \text{K}_{\text{BGMS}} \cdot \text{Revision}_{i,t} + \alpha + \text{u}_{i,t}$$

	(1)	(2)	(3)	(4)
	Unemployment		Inflation	
sample	1968-2017	1984-2017	1968-2017	1984-2017
$\text{Revision}_{i,t}$ (K_{BGMS})	0.321	0.398	0.143	-0.263
	(0.107)	(0.149)	(0.123)	(0.054)
R^2	0.028	0.052	0.005	0.025
Observations	5383	3769	5147	3643

Notes: The observation is a forecaster by quarter between Q4-1968 and Q4-2017. The forecast horizon is 3 quarters. Standard errors are clustered two-way by forecaster ID and time period. Both errors and revisions are winsorized over the sample to restrict to 4 times the inter-quartile range away from the median. The data used for outcomes are first-release ("vintage").

revisions for a particular forecaster, indexed by i. The regression of interest is the following:

$$\text{Error}_{i,t,k} = \alpha + K_{\text{BGMS}} \cdot \text{Revision}_{i,t,k} + u_{i,k,t} \tag{2}$$

where the main outcome and regressor are individual-level forecast errors and revisions; α is a constant; and K_{BGMS}, in shorthand reference to the authors of Bordalo et al. (2018), is the individual level analogue to K_{CG}.

Regardless of the information structure, individual-level rationality imposes $K_{BGMS} = 0$.[14] The literal interpretation of $K_{\text{BGMS}} > 0$ is that *individuals* update sluggishly in the direction of the truth; whereas if $K_{\text{BGMS}} < 0$, they generally revise beyond the realization of the data.

The fact: $K_{BGMS} < 0$ for inflation but $K_{BGMS} > 0$ for unemployment. In columns 1 and 3 of Table 3, we provide estimates of the individual-level regression (2) in the Survey of Professional Forecasters over the full sample. We focus on the horizon $k = 3$; results for other horizons are qualitatively similar and reported in Appendix Table A.5. Columns 2 and 4 of the same table conduct the analysis on the sub-sample from 1984 to the present.

For unemployment, we find substantial evidence that $K_{BGMS} > 0$ over the full and restricted sample period. And for inflation, we find imprecise evidence that $K_{\text{BGMS}} > 0$ over the full sample, including the 1970s and Volcker disinflation, but strong evidence of $K_{\text{BGMS}} < 0$ in the more stationary environment post 1984. The latter piece is consistent with the findings reported in Table 1, which showed that aggregate-level forecast inertia was much lower over this sample period.

These findings provide suggestive, if not completely conclusive, evidence that some departure from rational expectations is needed. Both Kohlhas and Broer (2019) and Bordalo et al. (2018) focus on $K_{\text{BGMS}} < 0$

[14] To be precise, $K_{BGMS} = 0$ follows from individual rationality together with perfect recall of one own's previous-period forecast revision. Imperfect recall could naturally give $K_{BGMS} > 0$ at the individual level, in a similar way that, as we will see more clearly in the next section, dispersed noisy information gives $K_{CG} > 0$. Intuitively, imperfect recall transforms a single forecaster to multiple "selves," one per period, each of whom has private information that can help forecast the forecast error of other "selves." As in the rest of the literature, we will largely ignore this possibility. But we will later explain why our own main lessons could be robust to it.

and argue that the necessary departure from rational expectations is a model in which agents over-react to recent information because they treat it as "over-representative" of the truth. Bordalo et al. (2018) offer a closely related explanation in terms of agents' being "over-confident" about the precision of their information.

In Section 5, we will illustrate how $K_{\text{BGMS}} < 0$ could indeed be explained by such over-confidence or over-representativeness. But we will also point out that another plausible departure from rational expectations, a form of over-extrapolation, could also give rise to $K_{\text{BGMS}} < 0$. This complicates the the structural interpretation of this moment of the expectations.

For now, we wish to iterate that while $K_{\text{BGMS}} < 0$ applies to forecasts of inflation (and also to forecasts of certain nominal interest rates, as shown in Bordalo et al., 2018), the opposite fact, $K_{\text{BGMS}} > 0$, applies to forecasts of unemployment. Therefore, the requisite departure from rational expectations is *not* the same across different forecasted variables.

The "right" variation. A final issue, that is ultimately unsurmountable in standard panel datasets, is that there is no analogue to our business-cycle-IV version of the CG regression at the individual level. It is hard to think of, let alone implement, an individual-specific treatment that isolates idiosyncratic information about a specific business-cycle shock (outside, potentially, a randomized control trial).[15] In that sense it is inevitable that the individual-level analysis conflate idiosyncratic interpretations of macro news and/or pieces of individual information about the macroeconomy, the theoretically "correct" variation, with slight differences in timing of response or actual mistakes filling out the form, a "less interesting" source of variation.

4 A New Fact: Delayed Over-shooting

So far we have focused on the unconditional relationship between forecast revisions and subsequent forecast errors. For our perspective, this approach has two main limitations. First, it "averages over" multiple shocks that may be driving actual outcomes and forecasts. This includes various "true" macroeconomic shocks of the kind included in mainstream models or identified in SVARs and more "residual" shocks, such as those associated with measurement error, correlated noise in information, or even seasonal fluctuations. And second, it gives a sense of the under-reaction of the average forecasts on impact but does not permit one to see the the dynamic adjustment of forecasts at longer horizons.

To address these limitations, we now focus on a pair of macroeconomic shocks and trace the IRFs of actual outcomes, forecasts, and forecast errors to these shocks.

[15] Recent contributions by Coibion, Gorodnichenko, and Kumar (2018) and Coibion, Gorodnichenko, and Ropele (2019) have this flavor, though the treatments do not necessarily embody the full general-equilibrium properties of a macro shock (i.e., the fact that everyone else will know of the change and also adjust their expectations, which will have real effects on outcomes).

4.1 Methodology

Identified shocks. As explained in more detail in Section 2, we consider two candidate shocks in the data that map roughly to "demand and supply shocks" in the theory (which we develop later on, in Sections 5 and 6). Both shocks are drawn from the SVAR analysis of Angeletos, Collard, and Dellas (2019).

The first shock is constructed by maximizing its contribution to the business cycle variation in unemployment and is found to have the following properties: it encapsulates strong positive co-movement in employment, output, investment, and consumption *only* over the business cycle; it has a negligible footprint on TFP at all horizons; and it has a small to modest footprint on inflation. It can thus be interpreted as an non- or mildly-inflationary demand shock.

The second shock is identified by maximizing its contribution to the business cycle variation in inflation and it is found to have a negative but very small footprint on real quantities and zero footprint on TFP. It can thus be interpreted as some kind of non-technology supply shock, or a markup shock, which manifests primarily in inflation.

One can quibble on whether these shocks are "truly" structural. This depends on what kind of models one has in mind. Angeletos, Collard, and Dellas (2019) show that the first of the aforementioned empirical shocks is closely related to certain theoretical counterparts: the investment-specific demand shock in Justiniano, Primiceri, and Tambalotti (2010), the risk shock in Christiano, Motto, and Rostagno (2014), and the confidence shock in Angeletos, Collard, and Dellas (2018). And the second of the two empirical shock is closely related to the markup or cost-push shock in such DSGE models.

Regardless of these specific structural interpretations, we contend that the aforementioned shocks are preferable for our purposes to other candidates, such as the technology and monetary shocks identified via timing restrictions in the SVAR literature (e.g., Galí, 1999; Sims and Zha, 2006), because they account for a much larger fraction of the business cycle variation in macroeconomic activity and inflation and may therefore also be the main driver in the corresponding variation in people's expectations.

We denote the two shocks, respectively, as $(\varepsilon_t^D, \varepsilon_t^S)$ for "demand" and "supply". We finally note that, while these shocks are not constructed to be orthogonal to one another, they are very close to being so in the data.

Main specification. To estimate dynamic responses to the aforementioned shocks, we consider two different empirical strategies.

The first is to estimate the IRFs via a parsimonious, instrumental-variables ARMA(P,K) representation. In particular, we estimate the following regression:

$$x_t = \alpha + \sum_{p=1}^{P} \gamma_p \cdot x_{t-p}^{\text{IV}} + \sum_{k=1}^{K} \beta_k \cdot \varepsilon_{t-k} + u_t \tag{3}$$

Depending on the variable whose dynamic response we want to look at, x_t is the actual outcome (unemployment or inflation), the relevant forecast, or the corresponding forecast error. In all cases, $\varepsilon_t \in \{\varepsilon_t^D, \varepsilon_t^S\}$ is one of the aforementioned two shocks drawn from Angeletos, Collard, and Dellas (2019). Finally, for

$p \in \{1, ..., P\}$, x_{t-p}^{IV} are the lagged values of x_t *instrumented* by the lagged values of ε_t.[16] This instrumental-variable approach allows us to recover the conditional dynamic responses to the structural shock under consideration—intuitively, how x_t moves when driven by the shock process of interest. We will call this method the "ARMA-IV" estimation.

By estimating (3) for outcomes (e.g., x_t equal to that quarter's unemployment rate or the past four quarters' inflation rate), we can generate dynamic impulse response coefficients $(\beta_{\text{out},h})_{h=0}^{H}$ as functions of $(\beta_0, (\gamma_p)_{p=1}^{P})$. For forecasts, we can do the same thing with x_t equal to the forecast in period t (e.g,. $\bar{\mathbb{E}}_t[u_{t+3}]$ and $\bar{\mathbb{E}}_t[\pi_{t+3,t-1}]$): estimate the impulse response coefficients $(\tilde{\beta}_{\text{fc},h})_{h=0}^{H}$ and then "re-index" these coefficients to line up with the realized outcomes. More specifically, we generate $(\beta_{\text{fc},h})_{h=0}^{H}$ such that $\beta_{\text{fc},h} = 0$ for $h < 3$ (effectively imposing unpredictability of the shocks), and $\beta_{\text{fc},h} = \tilde{\beta}_{\text{fc},h-3}$ for $h \geq 3$. Finally, we can construct the IRF of the forecast errors either by taking the difference between the IRF of the outcome and the forecasts, or by repeating the aforementioned procedure with x_t being the average forecast error.

In all cases, we construct standard errors for the coefficients that are heteroskedasticity and autocorrelation robust (HAC) with a 4-quarter Bartlett kernel; and then use the delta method to calculate standard errors for the impulse response functions. All reported error bands are 68% confidence intervals ($\pm 1 \cdot \text{SE}$).

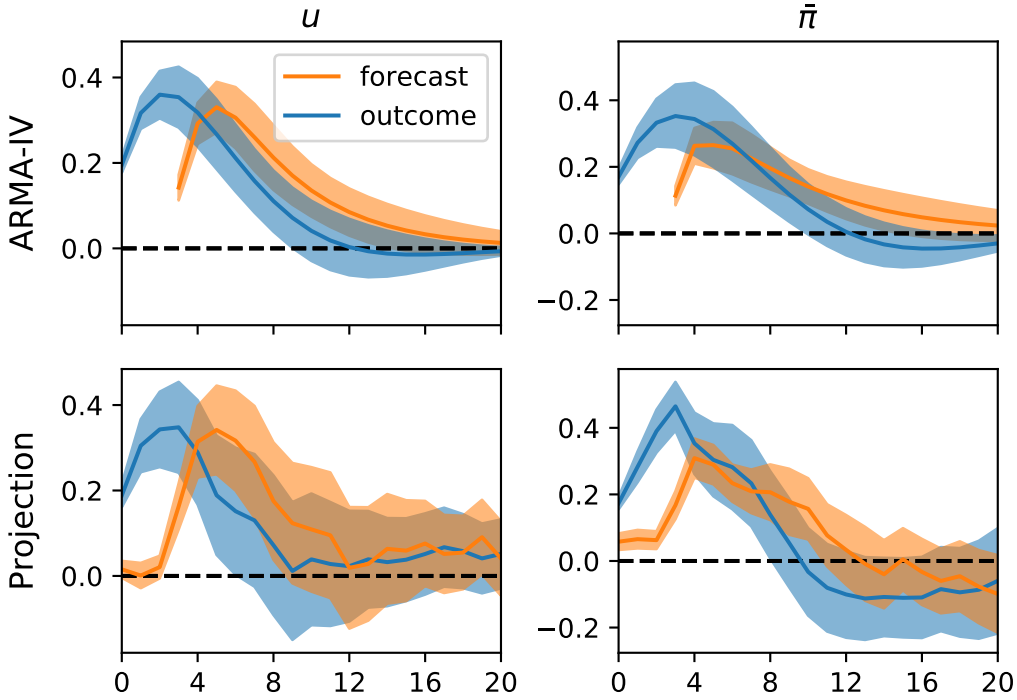
An "unrestricted" projection. Our main strategy strives for parsimony by requiring the IRFs to accept a low-dimension ARMA representation as in (3). But we can also estimate impulse responses directly using the projection method of Jordà (2005). In this case, the estimating equation, for each horizon $0 \leq h \leq H$, is the following:

$$x_{t+h} = \alpha_h + \beta_h \cdot \varepsilon_t + \gamma' W_t + u_{t+h} \tag{4}$$

where $(\beta_h)_{h=0}^{H}$ trace out the dynamic response of the outcome, W_t is a vector of control variables, and γ are the coefficients on these controls. Consistently, across specifications, we include the lagged outcome x_{t-1} and the lagged forecast $\bar{\mathbb{E}}_{t-k-1}[x_{t-1}]$ as control variables. Conceptually, as long as these controls are orthogonal to the structural shock ε_t, these should not affect the population estimate we get of the impulse response parameters; but their inclusion may help with small-sample precision. We find overall that results are not sensitive to choices of controls. Standard errors are constructed in the same, aforementioned way.

For our main analysis, we set $k = 3$ quarters as the forecast horizon and set $H = 20$ quarters as the maximum period for tracing out IRFs. The former assumption allows us to use the exact same data on unemployment and annualized inflation forecasts that formed the center of our analysis in Section 3.

Figure 1: Dynamic Responses in the Data: Outcomes and Forecasts



Notes: The sample period is Q1 1968 to Q4 2017. The shaded areas are 68% confidence intervals based on HAC standard errors with a Bartlett ("tent") kernel and 4 lags. The x-axis denotes quarters from the shock, starting at 0. In the first row the outcome is u_t and the forecast is $\bar{\mathbb{E}}_{t-3}[u_t]$; in the second row the outcome is $\pi_{t,t-4}$, or annual inflation, and the forecast is $\bar{\mathbb{E}}_{t-3}[\pi_{t,t-4}]$.

4.2 The fact: dynamic over-shooting

Figure 1 shows, in a two-by-two grid, the main impulse response estimates. In the first column, we show the dynamic response of unemployment and median forecasts thereof to the demand shock ε_t^D. The first row shows the instrumented ARMA method of equation (3) , and the second row shows the projection method of (4). For both methods, we "align" the forecast responses such that, at a given vertical slice of the plot, the outcome and forecast responses are measured over the same horizon, and the difference thereof is a measure of the response of forecast errors. In the second column, we plot the same for the response of one-year-average inflation to the supply shock ε_t^S.

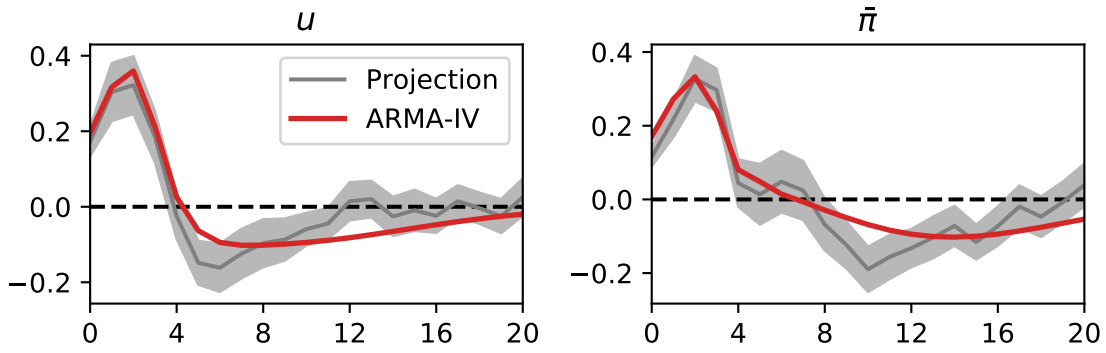
The consistent pattern across specifications is an initially delayed, and then over-persistent response of forecasts to the shock. Consider, as an illustration, the response of unemployment and forecasts thereof

[16]The first-stage equation is given, in vector form, by

$$X_{t-1} = \eta + \mathscr{E}'_{t-1}\Theta + e_t$$

where $X_{t-1} \equiv (x_{t-p})_{p=1}^{P}$, $\mathscr{E}_{t-1} \equiv (\varepsilon_{t-K-j})_{j=1}^{J}$ and $J \geq P$. Our main specifications use $P = 3$ and $J = 6$, but the results are robust to $P = 2$ and $P = 4$, as well as to different J.

Figure 2: Dynamic Responses in the Data: Forecast Errors



Notes: The sample period is Q1 1968 to Q4 2017. The shaded areas are 68% confidence intervals based on HAC standard errors with a Bartlett ("tent") kernel and 4 lags. The x-axis denotes quarters from the shock, starting at 0. In the first row the outcome is u_t and the forecast is $\bar{\mathbb{E}}_{t-3}[u_t]$; in the second row the outcome is $\pi_{t,t-4}$, or annual inflation, and the forecast is $\bar{\mathbb{E}}_{t-3}[\pi_{t,t-4}]$.

to ε_t^D. Unemployment spikes around quarter 3 in both estimation methods before reverting back to its long-run mean. The point-estimate is extremely close to zero by $t = 12$ in both cases.

Now consider the response of forecasts at $t = 3$ in the plot. These are forecasts *made* at $t = 0$, when the very first macro data (e.g., BLS reports) from $t = 0$ become available. Forecasted unemployment immediately spikes and begins to decay over the next 5-6 quarters. Forecasters remain convinced there are adverse demand conditions, when in reality conditions have reverted back to the mean. A similar, and indeed more dramatic, pattern is visible in the response of inflation to the supply shock (second row). And these patterns look qualitatively and quantitatively quite similar with both the smooth, ARMA estimates (left column) and the unrestricted projection regression estimates (right column).

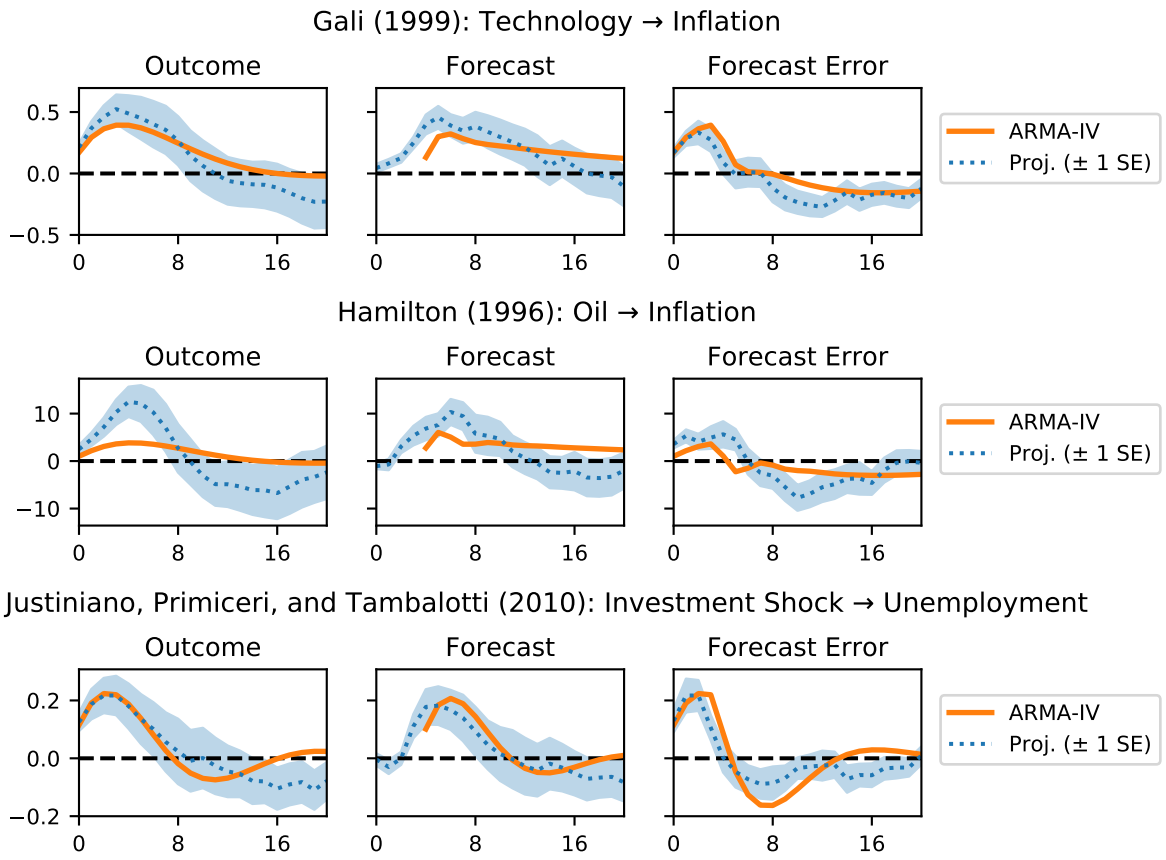
Figure 2 shows this overshooting pattern more clearly in terms of the impulse response of forecast errors. For both the ARMA and projection methods, this is obtained by taking the difference of the previous estimates for outcomes and forecasts. For both unemployment and inflation, we find evidence that forecast errors start positive and then turn negative at longer horizons. The estimated "crossing points" of the forecast errors response with 0, using the ARMA method, are $K^u_{\text{IRF}} = 4.14$ and $K^\pi_{\text{IRF}} = 6.43$, respectively.[17],

Finally, in Appendix Figure A.1, we complete the picture with the "off-diagonal" impulse responses of inflation to the demand shock and unemployment to the supply shock. The former is weakly inflationary at longer horizons and the latter weakly contractionary at medium horizons. And in both cases we have modest evidence of the over-shooting pattern of interest.

Other shocks of interest. An appealing feature of using projections and the ARMA-IV method is that it is easy to combine with auxiliary identification techniques, without fully specifying a multivariate model and considering the problem of jointly identifying many shocks. To illustrate this property, and probe the robustness of our results to other candidate "supply and demand" shocks from the macroeconomics

[17]The corresponding estimates from the projection regressions are 4.87 and 7.79.

Figure 3: Responses to Other Structural Shocks

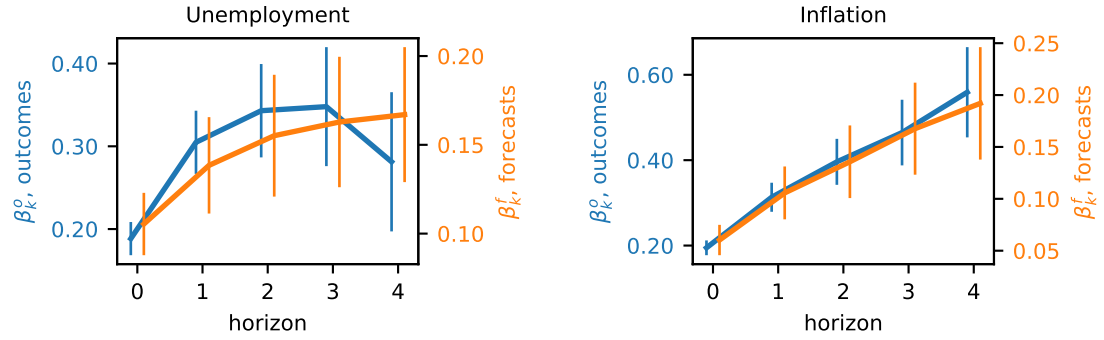


Notes: The sample period is Q4 1968 to Q4 2017. The x-axis denotes quarters from the shock (starting at 0). The shaded areas are 68% confidence intervals based on HAC standard errors with a Bartlett ("tent") kernel and 4 lags. The x-axis denotes quarters from the shock, starting at 0. The first shock is a technology shock à la (Galí, 1999), as obtained from Coibion and Gorodnichenko (2012) and normalized to be inflationary and contractionary. The second is an oil shock à la Hamilton (1996), again obtained from Coibion and Gorodnichenko (2012). The third is the investment-specific shock of Justiniano, Primiceri, and Tambalotti (2010), updated to cover the full sample until 2017. See Appendix B for details.

literature, Figure 3 replicates our main analysis for three different shocks: a technology shock à la Galí (1999), normalized here to be inflationary and contractionary; an oil price shock à la Hamilton (1996); and the investment-specific shock extracted from the DSGE model of Justiniano, Primiceri, and Tambalotti (2010). The former two are variations of "supply shocks" (to productivity or input costs), and we show the response of inflation; the last is like a demand shock, and we show the response of unemployment. In all cases we see evidence of the overshooting pattern.

Methods for estimating dynamics. Our ARMA-IV method resembles the technique suggested by Romer and Romer (2004), and applied by Coibion and Gorodnichenko (2012) in their study of how forecast errors respond to structural shocks. The Romer and Romer (2004) technique estimates an empirical ARMA process like (3) via ordinary least squares ("ARMA-OLS"). This method uses the *unconditional* auto-covariance

Figure 4: The Term Structure of Forecasts and Outcomes



Notes: The sample period is Q4 1968 to Q4 2017. The x-axis denotes quarters from the shock or horizon of forecast (starting at 0). The lines are one-standard-error bars. The orange lines plot the terms structure of forecasts, or β^f_k from (5), and the blue lines show the response of outcomes, or β^o_k from (6).

properties of the outcome variable in order to quantify dynamics. Our prior is that, in a world of very different, shock-specific dynamics (induced, for instance, by differential persistence in the driving process or differential ability to learn about these shocks), the ARMA-OLS method could give mis-leading results. Indeed, in our replication of two key results from Coibion and Gorodnichenko (2012), the response of inflation and forecast errors thereof to technology and oil shocks (Figure 3), we find evidence of our overshooting patterns when we use our methods. Appendix B unpacks more thoroughly the differences in methodology and demonstrates that both using unconditional dynamics and restricting, via model-selection tools, to models with only one root in dynamics makes it impossible to see the over-shooting patterns uncovered here.

Another option for estimating more complex dynamics, of course, is to jointly estimate a multivariate model. A VAR model will use unconditional information (i.e., the reduced-form representation) to project forward dynamics of the identified shocks—but it will have much more richness in modeling cross-variable interactions, and will certainly allow two shocks with equal instantaneous effects on a variable of interest (say, unemployment) but different effects on other variables (say, labor productivity or TFP) to have different dynamic responses. We will pursue such a strategy in the next section.

The "term structure" of forecasts. Finally, a complementary way of organizing the evidence on overshooting is to focus on the "term structure" of forecasts at a given point in time—how much do professionals move their long-horizon forecasts (say, one-year-out) versus their now-casts? To estimate this in the data, we consider the following "slice" of the projection regression for forecasted variable x at different horizons k:

$$\bar{\mathbb{E}}_t[x_{t+k}] = \alpha_k + \beta^f_k \cdot \varepsilon_t + \gamma' W_t + u_{t+k} \tag{5}$$

and the same for realized outcomes

$$x_{t+k} = \alpha_k + \beta_k^o \cdot \varepsilon_t + \gamma' W_t + u_{t+k} \tag{6}$$

We run these specifications for x equal to unemployment and (cumulative k-period) inflation. For consistency, we use the two control variables corresponding to horizon $k = 3$ that we used in the projection regression (4)—that is, u_{t-1} and $\bar{\mathbb{E}}_{t-4}[u_{t-1}]$ for unemployment and $\pi_{t-1,t-5}$ and $\bar{\mathbb{E}}_{t-4}[\pi_{t-1,t-5}]$ for inflation.

The coefficients of interest are (β_k^f, β_k^o), which reveal the persistence of outcomes and forecasts. If $\beta_k^f < \beta_k^o$, which we have already verified for $k = 3$, we know agents under-react on impact. If β_k^o is much more persistent across k than β_k^f, this is also evidence of over-extrapolation right at the impact of the shock—that is, agents end up being *more correct* about impacts further in the future because their over-extrapolation partially cancels out their under-reaction.

Figure 4 plots the results, showing the sequences of (β_k^o, β_k^f) on the left and right scales, respectively, for unemployment (left graph) and inflation (right graph). Observe both that agents under-react (comparing the left and right scales), but also that under-reaction is most severe at the two *shorter* horizons. At horizon 0, agents incorporate into their forecasts $R_0 \equiv 100 \cdot \beta_0^f / \beta_0^o = 100 \cdot 0.105/0.188 = 55.9\%$ of the forecast, compared with 45.4% and 59.4% respectively at horizons 1 and 4. Comparing $k = 4$ to $k \in \{0, 1\}$, the term structure "slopes up"—agents are actually more correct about the further future. For inflation, the corresponding "ratios" at horizons 0, 1, and 4 are 31.0%, 33.8%, and 34.4%—fairly flat but still upward sloping.

4.3 A Structural VAR Approch

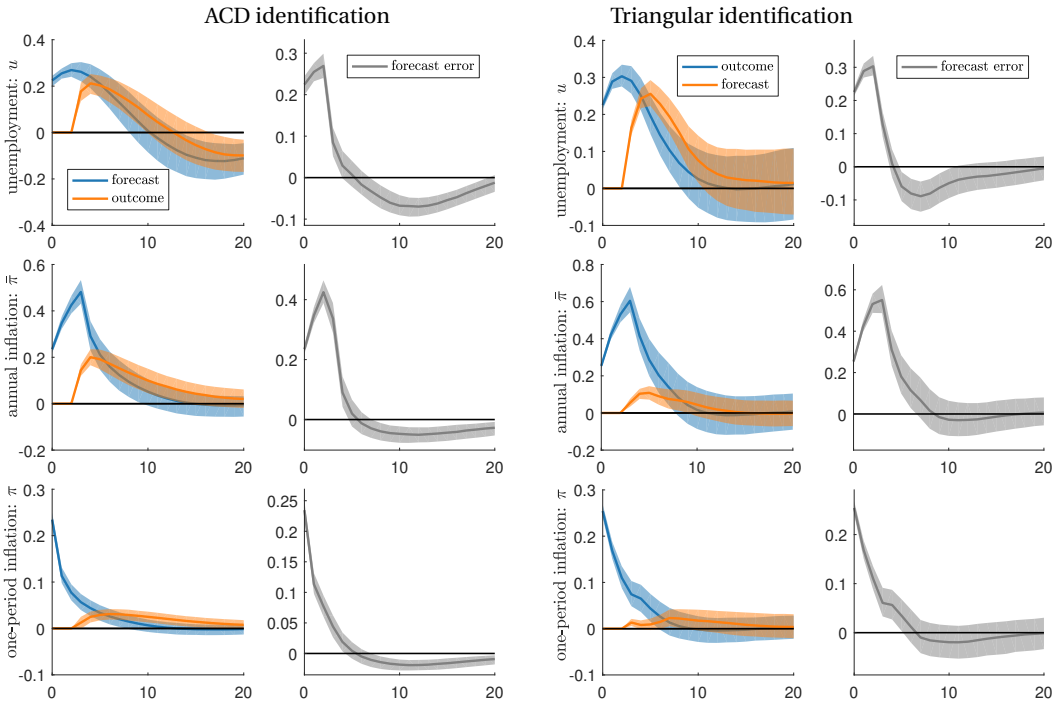
We now probe the robustness of our result to estimation in a multi-variate VAR model.

We consider two specific models. Both are based on the same 14-variable reduced-form VAR comprised of the ten key macroeconomic variables from Angeletos, Collard, and Dellas (2019) plus the three forecast variables: the three-period-ahead unemployment forecast, the three-period-ahead annual inflation forecast, and the three-period-ahead quarterly inflation forecast. The macro variables are the following: real GDP, real investment, real consumption, labor hours, the labor share, the Federal Funds Rate, labor productivity, and utilization-adjusted TFP.[18] The forecast variables are are the three-step-ahead unemployment and inflation forecasts from the SPF. The sample period is Q4 1968 to Q4 2017. We apply the same Bayesian inference procedure as Angeletos, Collard, and Dellas (2019), including prior specification and posterior sampling procedures, and replicate their identification of shocks that target the "max share of variation" in both unemployment and inflation.

On the left side of Figure 5, we show what happens when we replicate the identification scheme of Angeletos, Collard, and Dellas (2019). In the first row, we show the response of unemployment, forecasts thereof, and forecast errors to the "unemployment shock." This can be compared directly to the first column of Figures 1 and 2, and largely agrees about the potential for large and persistent "overshooting" in forecast errors. The second and third row show the response of outcomes and forecasts to the inflation

[18] Full variable descriptions and data construction discussion is in Angeletos, Collard, and Dellas (2019).

Figure 5: Dynamic Responses in a Structural VAR



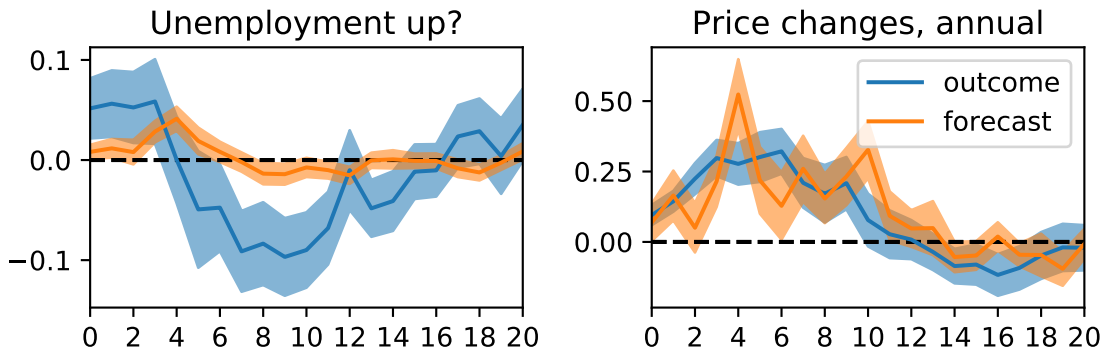
Notes: The sample period is Q4 1968 to Q4 2017. The x-axis denotes quarters from the shock (starting at 0). The shaded areas are 68% high-posterior-density regions and the point estimate is the posterior median. In the first row the outcome is u_t and the forecast is $\bar{\mathbb{E}}_{t-3}[u_t]$; in the second row the outcome is $\pi_{t,t-4}$, or annual inflation, and the forecast is $\bar{\mathbb{E}}_{t-3}[\pi_{t,t-4}]$; and in the last row, the outcome is $\pi_{t,t-1}$, or one-quarter inflation, and the forecast is $\bar{\mathbb{E}}_{t-3}[\pi_{t,t-1}]$. The columns show results from a "max share" identification and a triangular identification, respectively; see the main text for details.

shock in the same SVAR model, but with different forecast horizons and transformations of the outcome variable (annual averages in Row 2 versus quarter-to-quarter rates in Row 3). Here we find quantiatively smaller effects per period, but also very persistent ones. In Appendix Figure A.2, we show the "missing" impulse responses of unemployment to the supply shock and inflation to the demand shock; they, too, match the patterns in the projections of Figure A.1 and show evidence of the overshooting.

On the right side of Figure 5, we show just the results of two different "Cholesky" identifications based on triangular short-run restrictions (ordering unemployment or inflation first). The first row shows the impulse response of unemployment, forecasts thereof, and forecast errors to the "one-step-ahead" shock to unemployment (i.e., the shock in the triangular-identified SVAR in which unemployment is the "slowest" variable, hit only by one shock). We find corroborating evidence of the "over-shooting" pattern documented earlier. The second row and third rows show, respectively, the response of inflation and forecasts thereof in either the quarter-to-quarter or annual-average units to the one-step-ahead shock to inflation (i.e., when inflation is ordered first). Here, there is evidence of over-shooting in point estimate, but not very large magnitudes or precision.

Taking the evidence together, we conclude that the fact for unemployment is particularly robust to

Figure 6: Dynamic Responses in the Michigan Survey



Notes: The sample period is Q1 1968 to Q4 2017. The shaded areas are 68% confidence intervals based on HAC standard errors with a Bartlett ("tent") kernel and 4 lags. In the first plot the outcome is $\text{UnempUp}_{t,t-4}$ and the forecast is $\bar{\mathbb{E}}_{t-4}[\text{UnempUp}_{t,t-4}]$; in the second plot the outcome is $\pi_{t,t-4}$, or annual inflation, and the forecast is $\bar{\mathbb{E}}_{t-4}[\pi_{t,t-4}]$.

different shock identification methods. The result for inflation is somewhat more dependent on the particular variation on which we focus.

4.4 "General public" forecasts

We also investigate whether these patterns hold in *non-professional* forecast data. Our source for these data is the University of Michigan Survey of Consumers. We construct an "unemployment expectation" using the survey's question about whether unemployment will go up, stay the same, or go down over the next 12 months. We code a variable $\bar{\mathbb{E}}_t[\text{UnempUp}_{t+4}]$ that averages the "up" responses, and code a data equivalent UnempUp_{t+4} using the BEA unemployment rate.[19] For inflation, we use the Michigan survey's main estimate for inflation over the next 12 months. For consistency with the previous analysis, we compare this to data on the GDP deflator, even though this is almost certainly not a perfect match for the price variable households have in mind when making their forecast.[20]

Figure 6 shows the results from projecting our business cycle shocks on these variables using (4). The left panel shows the response of the UnempUp variable and forecasts thereof to ε_t^D. The Michigan survey expectations perk up slightly *before* the shock hits (i.e.,. for $t < 4$) and then spike one quarter "too late." We see further evidence that the general public is also particularly unable to forecast the "mean-reverting" part of the shock, or the eventual downward trend in unemployment.

The right panel shows the response of the response of GDP deflator and the annual inflation expectation of the Michigan survey to ε_t^S. Here, responses are much too noisy to pick out an obvious "peak response." Again, there is some weak evidence of anticipation, and at quarters 10 and onward evidence of

[19] Results are similar if we treat a different portion (e.g., 1/2 or all) of the "about the same" responses as corresponding to "up."

[20] See the exact question in Footnote 7, which refers to "prices" in general and may match some notion of a consumer spending basket better than the GDP deflator basket.

some over-extrapolation of recent price trends.

4.5 Robustness: different data and methodologies

We conduct a number of additional robustness checks and extensions, the results of which are reported in the Appendix.

Appendix Figure A.3 replicates all of the previous patterns using the "vintage" data on unemployment and GDP/GNP deflator that are used in the forecast error predictability exercises, and finds consistent patterns. Appendix Figure A.4 recreates the regression results in the SPF, back again with modern-vintage data, in the sample 1984-2017. As discussed previously, we might think of the post-Volcker and post-oil-crisis data as a more "consistently stationary" regime for forecasters trying to model the world. We find largely the same patterns in forecast errors. Appendix Figure A.5 recreates the main analysis with forecast data from Blue Chip Economic Indicators over the shorter available sample (1980-2017) and again finds the same patterns.

5 A Simple Model

The overarching goal of the remainder of this article is to identify a parsimonious model that explains the previous three facts, and then use this model to determine the "causal effect of imperfect expectations"—that is, compare the real world with the counterfactual world of "perfect rational expectations." The latter task will involve fully accounting for general equilibrium feedback, or the endogeneity of the forecasted outcomes to forecasts. We grapple with this task fully in Sections 6 and 7. This section instead concentrates on the intermediate step of differentiating mechanisms for imperfect expectations in a much simplified environment that treats the forecasted outcomes as exogenous. A principal advantage of this approach is that it is simpler, and in many cases produces closed-form mappings between theoretical forces of interest and the moments we have measured so far.

Our main finding is that a combination of *noise* and *over-extrapolation* is needed to match Facts 1-3. On the way, we also explain how to "correct" the original CG coefficient so as to partial out the effect of irrationality and isolate the effect of noise.

5.1 Primitives

Let $\{z_t\}$ be a stochastic process that a group of agents, indexed by $i \in [0,1]$, are trying to forecast. Ultimately, we want to think of these processes as endogenous to the agents' behavior. But for now we assume that z_t follows an exogenous first-order auto-regressive process with Gaussian errors. Let $\rho \in (0,1)$ be the persistence parameter and $\varepsilon_t \sim N(0,1)$ a Gaussian innovation. The evolution of $\{z_t\}$ can be described as

$$(1-\rho\mathbb{L})z_t = \varepsilon_t, \tag{7}$$

where $\mathbb{L}$ denotes the lag operator (i.e., $\mathbb{L}z_t = z_{t-1}$).

Agents do not observe z_t, but instead only private signals of it. An agent's signal is contaminated with idiosyncratic noise $\tau^{-1/2}u_{i,t}$, where τ is a precision parameter and $u_{i,t} \sim N(0,1)$ is idiosyncratic Gaussian noise. In math, the actual realization of $s_{i,t}$ is given by

$$s_{i,t} = z_t + \frac{u_{i,t}}{\sqrt{\tau}} \tag{8}$$

As in a large literature, we can think of this noise either literally, as the product of dispersed noisy information (Lucas, 1972; Morris and Shin, 2002; Lorenzoni, 2009), or metaphorically, as a representation of rational inattention and bounded information processing capacities (Sims, 2003, 2010; Woodford, 2003, 2009; Mackowiak and Wiederholt, 2009; Mankiw and Reis, 2002). Unlike these earlier works, however, we combine "rational confusion" with two forms of irrationality.

First, we allow agents to have a misspecified belief about the process of the object they are trying to forecast. Whereas the true process from z_t is given by (7), agents perceive this process to be

$$(1 - \hat{\rho}\mathbb{L})z_t = r\varepsilon_t \tag{9}$$

for some *perceived* persistence $\hat{\rho}$ which may not equal ρ.

And second, we allow agents to have a misspecified belief about their information. Whereas the *true* process of the private signal is given by (8), agents perceive this process to be

$$s_{i,t} = z_t + \frac{u_{i,t}}{\sqrt{\hat{\tau}}} \tag{10}$$

for some *perceived* precision $\hat{\tau} > 0$ that may differ from τ.

We can think of $\hat{\tau} > \tau$ as a model of individual overconfidence: each forecaster systematically thinks their information is better than it truly is. The opposite case of $\hat{\tau} < \tau$ captures individual underconfidence: an individual thinks their information is systematically *worse* than it really is. Moore and Healy (2008) provide a representative review of the experimental psychological evidence for such biases. Their broad conclusion is that overconfidence is consistently prevalent for reported beliefs in the laboratory, but that the extent of effects can be context-specific.

The case $\hat{\rho} > \rho$ encodes an "over-extrapolation" of today's state to tomorrow, whereas $\hat{\rho} < \rho$ encodes "under-extrapolation." Both narratives are appealing in different economic contexts. On the one hand, Greenwood and Shleifer (2014) and Gennaioli, Ma, and Shleifer (2015a) argue that over-extrapolation, or $\hat{\rho} > \rho$, is evident both in stock-market expectations and in expectations of firms' sales forecasts; see also Guo and Wachter (2019) for how a simple model with over-extrapolation over dividend growth can explain a variety of asset-price phenomena. On the other hand, level-K thinking (Garcıa-Schmidt and Woodford, 2019; Farhi and Werning, 2019) and cognitive discounting (Gabaix, 2019) are close "close in spirit" to the opposite scenario, $\hat{\rho} < \rho$, because they cause agents to be under-estimate the (endogenous or exogenous) response of future outcomes to current innovations. We will make this connection formal in Section 6.4.3, once we extend the analysis to a GE context and properly nest these models.

5.2 Characterization of forecasts and errors

The agent's information set in period t is given by $\mathcal{I}_{i,t} = \{s_{i,t-v}\}_{v=0,1,...}$, the history of these signals up to period t. We will use the short-hand notation $\mathbb{E}_{i,t}[\cdot]$ to denote the agent's *subjective* expectation conditional on this information.

Our main goal in this section is to construct the theoretical counterparts of the various forecast moments we estimated in the empirical part of the paper. To this goal, let us first note that the law of motion for $\mathbb{E}_{i,t}[z_t]$, the "nowcast," can then be written as follows:

$$\mathbb{E}_{i,t}[z_t] = (1-\hat{g})\mathbb{E}_{i,t-1}[z_t] + \hat{g}s_{i,t} = (1-\hat{\lambda})\mathbb{E}_{i,t-1}[z_{t-1}] + \left(1-\frac{\hat{\lambda}}{\hat{\rho}}\right)s_{i,t}$$

where

$$\hat{g} \equiv 1-\frac{\hat{\lambda}}{\hat{\rho}} \in (0,1)$$

is the Kalman gain and $\hat{\lambda}$ is the unique root within $(0,\hat{\rho})$ to the following quadratic equation:

$$\hat{\lambda} + \frac{1}{\hat{\lambda}} = \hat{\rho} + \frac{1+\hat{\tau}}{\hat{\rho}} \tag{11}$$

All this is exactly as in the textbook Kalman filter, except for the fact that $\hat{\rho}$ and $\hat{\tau}$ have taken the place of the corresponding true parameters, since we are describing the evolution of *subjective* expectations.

Using the above result, we can characterize the k-step ahead forecasts and the relevant forecast errors and forecast revisions as follows:

Lemma 1. *The one-step-ahead forecasts obey*

$$\mathbb{E}_{i,t}[z_{t+1}] = (\hat{\rho}-\hat{\lambda})\frac{1}{1-\hat{\lambda}\mathbb{L}}s_{i,t} = (\hat{\rho}-\hat{\lambda})\frac{1}{1-\hat{\lambda}\mathbb{L}}\left(\frac{1}{1-\rho\mathbb{L}}\varepsilon_t + u_{i,t}\right) \tag{12}$$

The corresponding forecast errors obey

$$Error_{i,t} \equiv z_{t+1} - \mathbb{E}_{i,t}[z_{t+1}] = \frac{1-\hat{\rho}\mathbb{L}}{(1-\rho\mathbb{L})(1-\hat{\lambda}\mathbb{L})}\varepsilon_{t+1} - (\hat{\rho}-\hat{\lambda})\frac{1}{1-\hat{\lambda}\mathbb{L}}u_{i,t} \tag{13}$$

And finally the forecast revisions obey

$$Revision_{i,t} \equiv \mathbb{E}_{i,t}[z_{t+1}] - \mathbb{E}_{i,t-1}[z_{t+1}] = \left(\hat{\rho}-\hat{\lambda}\right)\frac{1}{1-\hat{\lambda}L}\left(\frac{1-\hat{\rho}L}{1-\rho L}\varepsilon_t + u_{it} - \hat{\rho}u_{it-1}\right) \tag{14}$$

Three properties are worth noting. First, the parameter controlling the persistence of forecasts, the previously defined $\hat{\lambda}$, decreases in perceived signal precision $\hat{\tau}$, to a minimum value of 0, from a maximum value of $\hat{\rho}$. Second, the forecast error 13 contains both aggregate terms (i.e., functions of $(\varepsilon_s)_{s<t}$) and idiosyncratic terms (i.e., functions of $(u_{i,s})_{s<t}$).

5.3 Forecast error predictability (Facts 1 and 2)

We are now ready to offer a structural interpretation to Facts 1 and 2, regarding the predictability of forecast errors by lagged forecast revisions.

Proposition 1 (Facts 1 and 2)**.** *Let* $\mathcal{V}_{ind} \equiv Var[\mathbb{E}_{i,t}[x_{t+1}] - \mathbb{E}_{i,t-1}[x_{t+1}]]$ *and* $\mathcal{V}_{agg} \equiv Var[\bar{\mathbb{E}}_t[x_{t+1}] - \bar{\mathbb{E}}_{t-1}[x_{t+1}]]$ *be the variances of, respectively, individual and aggregate forecast revisions; the former is a function of* $(\hat{\tau}, \rho, \hat{\rho})$ *and* τ*, the latter is a function of* $(\hat{\tau}, \rho, \hat{\rho})$ *but not of* τ*. The theoretical counterparts of coefficients of regressions* (2) *and* (1) *are given by, respectively,*

$$K_{BGMS} = \frac{-\kappa_1(\tau^{-1} - \hat{\tau}^{-1}) + \kappa_2(\rho - \hat{\rho})}{\mathcal{V}_{ind}} \tag{15}$$

$$K_{CG} = \frac{\kappa_1 \tau^{-1}}{\mathcal{V}_{agg}} + \frac{\mathcal{V}_{ind}}{\mathcal{V}_{agg}} K_{BGMS} \tag{16}$$

for some $\kappa_1 > 0$ *and* $\kappa_2 > 0$ *that are functions of* $(\hat{\tau}, \rho, \hat{\rho})$ *but not on* τ*.*

Condition (15) illustrates how the sign of K_{BGMS}, which pertains to the individual-level regression (2), hinges on the two departures from rational expectations. Under rational expectations (which herein means $\hat{\tau} = \tau$ and $\rho = \hat{\rho}$), we have $K_{\text{BGMS}} = 0$. Relative to this benchmark, both overconfidence ($\hat{\tau} > \tau$) and over-extrapolation ($\hat{\rho} > \rho$) contribute towards $K_{\text{BGMS}} < 0$. And the converse is true for underconfidence or under-extrapolation.

Corollary 1. *The following properties hold for* K_{BGMS}*:*

1. $\hat{\tau} = \tau$ *and* $\hat{\rho} = \rho$ *(noisy but rational expectations) restrict* $K_{BGMS} = 0$*.*
2. $\hat{\tau} \geq \tau$ *(overconfidence) and* $\hat{\rho} \geq \rho$ *(over-extrapolation), with complementary slackness, imply* $K_{BGMS} < 0$*, or over-reaction of individual forecasts in the sense of regression* (2)*.*
3. $\hat{\tau} \leq \tau$ *(underconfidence) and* $\hat{\rho} \leq \rho$ *(under-extrapolation), with complementary slackness, imply* $K_{BGMS} > 0$*, or over-reaction of individual forecasts in the sense of regression* (2)*.*

Condition (16) shifts focus to the aggregate-level regression (1) and the sign of K_{CG}. To understand this condition, let us first focus on the special case of rational expectations (which herein means $\hat{\tau} = \tau$ and $\rho = \hat{\rho}$). The formula for the CG regression coefficient then reduces to

$$K_{\text{CG}} = \frac{\kappa_1 \tau^{-1}}{\mathcal{V}_{\text{agg}}}.$$

If we use the formula for κ_1 (which can be found in the Appendix) and calculate $\mathcal{V}_{\text{agg}}$, we can re-express the above more simply as follows:

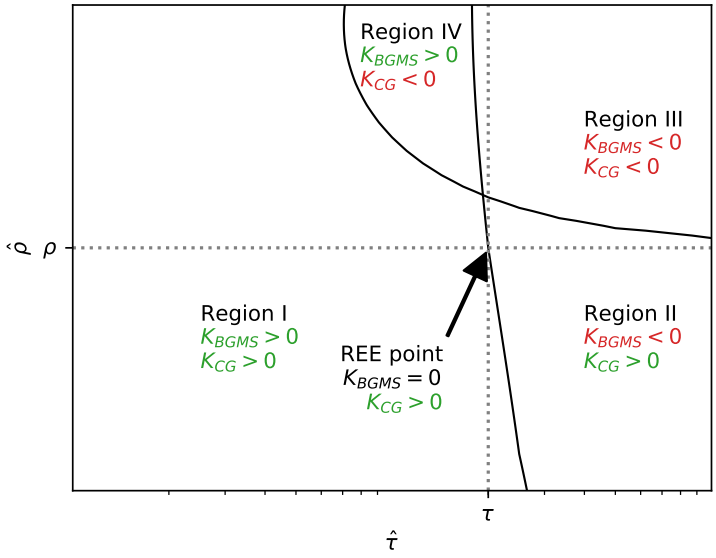
$$K_{\text{CG}} = \frac{1-g}{g}$$

where $g \in (0,1)$ is now the *objective* Kalman gain.[21] It follows that K_{CG} is decreasing in τ, or equivalently increasing in the level of noise, with $K_{\text{CG}} = 0$ when the noise is zero and $K_{\text{CG}} \to \infty$ as the noise becomes infinite.

This sums up the structural interpretation adopted in Coibion and Gorodnichenko (2015): in their framework, the coefficient K_{CG} obtained from regression (1) is interpreted as a direct measure of the informational friction. The *essence* of this structural interpretation extends to the larger class of models that

[21] Although $K_{\text{CG}} = \frac{1-g}{g}$ when $\tau = \hat{\tau}$ and $\hat{\rho} = \rho$, in general $K_{\text{CG}} \neq \frac{1-\hat{g}}{\hat{g}}$.

Figure 7: Parameter Cases for K_{BGMS} and K_{CG} in the Theory



allow for various kinds of informational frictions and a GE feedback between expectations and outcomes, while maintaining the pillar of rational expectations.[22] In this class of models, the exact mapping from K_{CG} to the primitive informational parameter is more convoluted, but K_{CG} can differ from zero *only* because of the informational friction. In this sense, K_{CG} alone remains a useful measure of the informational friction.

Our result illustrates how this structural interpretation crucially depends on ruling out a departure from rationality. When, instead, agents have a misspecified model for either their information ($\hat{\tau} \neq \tau$) or the stochastic process of the object they are trying to forecast ($\hat{\rho} \neq \rho$), the CG coefficient *confounds* the informational friction with the departure from rationality. In particular, $K_{\text{CG}} > 0$ could mean either that there is an informational friction or that there is a departure from rationality in the particular direction of underconfidence or under-extrapolation.

We summarize these lessons below.

Corollary 2. *The following properties hold for K_{CG}:*

1. *When $\hat{\tau} = \tau$ and $\hat{\rho} = \rho$ (noisy but rational expectations), K_{CG} is non-negative and strictly increasing in the level of noise. In this sense, K_{CG} is a measure of the informational friction.*

2. *More generally, the measure of the informational friction contained in K_{CG} is contaminated by the departure from rational expectations: a high positive value for K_{CG} could also be the implication of underconfidence ($\hat{\tau} < \tau$) or under-extrapolation ($\hat{\rho} < \rho$).*

[22] This includes models with sticky information (Mankiw and Reis, 2002), rational inattention (Sims, 2003; Mackowiak and Wiederholt, 2009), endogenous signals (Lucas, 1972; Nimark, 2008), and rich higher-order uncertainty (Acharya, Benhabib, and Huo, 2017; Angeletos and La'O, 2013; Bergemann and Morris, 2013; Nimark, 2008).

Facts 1 and 2 combined. Figure 7 illustrates how the combination of K_{CG} and K_{BGMS} depends on the combination of the two forms of mis-specification allowed here. The figure presumes a certain pair of values for ρ and τ, and shows how the space of $\hat{\rho}$ and $\hat{\tau}$ can be split in four areas, each corresponding to a different combination of signs for the CG and BGMS regression coefficients.

In Region I, which corresponds to large enough underconfidence and/or large enough under-extrapolation, both coefficients are positive (and in this sense we forecasts under-react at both the individual and the aggregate level). In Region III, which corresponds to large enough overconfidence and/or large enough over-extrapolation, we get the exact opposite: both coefficients to are negative (or forecasts over-react at both the individual and the aggregate level). Finally, in Regions II and IV, the two coefficients have opposite signs (or forecasts exhibit under-reaction in one level and over-reaction in the other).

By combining this figure with Facts 1 and 2 from the previous section, we can indeed infer that the empirically relevant case is a point within Region I for the case of unemployment forecasts and a point within Region II for the case of inflation forecasts. But the construction of this figure presumes knowledge of both ρ and τ. The analyst ("econometrician") may be able to identify ρ from the time series of the actual outcome (unemployment or inflation, depending on the case considered). But how can she possible identify τ, or the true level of noise?[23]

To identify all the three belief-related parameters, namely τ, $\hat{\tau}$, and $\hat{\rho}$, we need three moments. The two moments are those already discussed: the empirical estimates of K_{CG} and K_{BGMS} from Facts 1 and 2. The third relates to Fact 3.

5.4 Dynamic responses: under-shooting early, over-shooting later on (Fact 3)

Let us now turn to the structural interpretation of Fact 3, the sign reversal in the IRF of the aggregate forecast errors. This will complete the explanation of how the three facts fit in our framework, and the foundation of our identification strategy.

Proposition 2 (Fact 3)**.** *Let $\{\zeta_k\}_{k=1}^{\infty}$ be the Impulse Response Function (IRF) of the average forecast error. That is, for all $k \geq 1$,*

$$\zeta_k \equiv \frac{\partial\left(z_{t+k} - \overline{\mathbb{E}}_{t+k-1}[z_{t+k}]\right)}{\partial \varepsilon_t}$$

is the k-th coefficient in the moving-average representation of the average forecast error.[24]

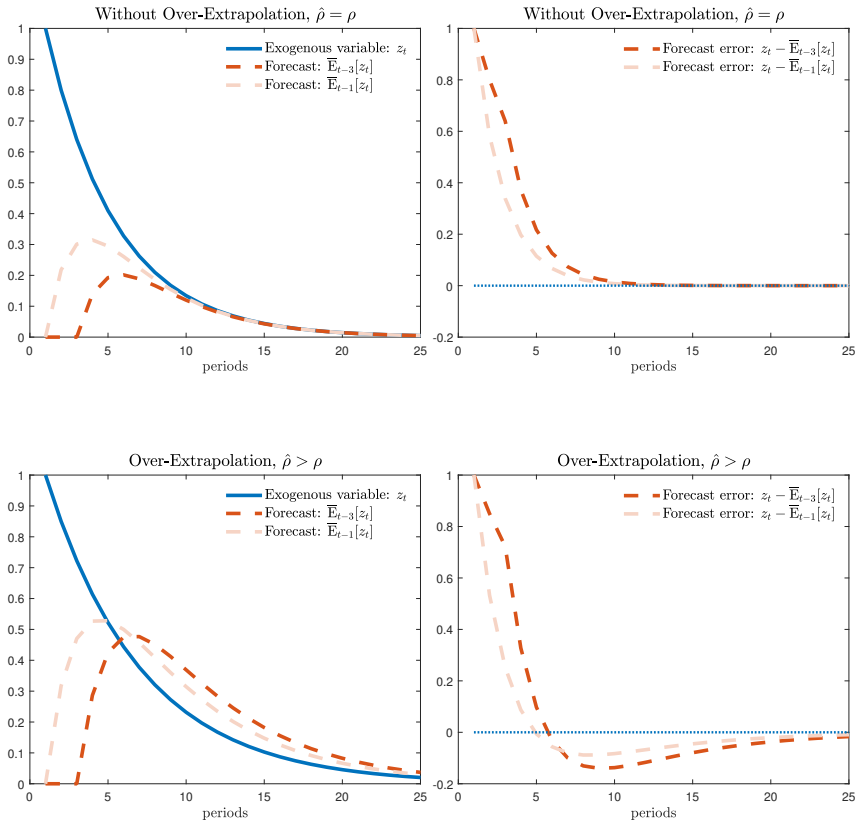
1. *If $\hat{\rho} < \rho$, or agents under-extrapolate, then $\zeta_k > 0$ for all $k \geq 1$. That is, the IRF of the average forecast error is uniformly positive.*

2. *If $\hat{\rho} > \rho$ and $\hat{\lambda} > \hat{\rho} - \rho$, or agents over-extrapolate and learning is slow enough, then $\zeta_k > 0$ for $1 \leq k <$*

[23] It is true of course that Facts 1 and 2 together say something about this level of noise, albeit not independently from other parameters. Appendix C shows more clearly how to make this connection and use a hybrid regression of individual and aggregate predictability to test for noisy signals in this class of models.

[24] Note that we exclude $\zeta_0 \equiv \partial\left(z_t - \overline{\mathbb{E}}_{t-1}[z_t]\right)/\partial \varepsilon_t$ for the present statement, because this is mechanically 1 no matter the belief structure.

Figure 8: IRF of Aggregate Forecasts and Errors in the Theory (without GE)



K_{IRF} and $\zeta_k < 0$ for $k > K_{IRF}$, where

$$K_{IRF} = \frac{\log(\hat{\rho} - \rho) - \log(\hat{\rho} - \hat{\lambda})}{\log \hat{\lambda} - \log \rho} > 1 \tag{17}$$

That is, the IRF of the average forecast errors starts positive but eventually switches negative.

3. *Finally, if $\hat{\rho} > \rho$ but $\hat{\lambda} < \hat{\rho} - \rho$, or agents over-extrapolate but learning is fast, then $\zeta_k < 0$ for all $k \geq 1$. That is, the IRF of the average forecast errors is uniformly negative.*

There are two key take-aways. First, a "sign-switch" in the impulse response of forecast errors to a macro shock necessitates the *combination* of over-extrapolation and noise. Second, the speed with which the sign flip occurs (e.g., how small is K_{IRF}) provides direct evidence on the extent of over-extrapolation relative to the informational friction.[25]

Figure 8 illustrates these patterns by plotting the IRFs of outcomes and forecasts (left column) and forecast errors (right column) in two scenarios: a benchmark without over-extrapolation (top row), and a variant with (bottom row). In each case, we report both the one-step-ahead forecasts and errors (light dashed lines) and their three-step-ahead counterparts (dark dashed lines). The former correspond the

[25] On a more technical level, note that, as written, K_{IRF} need not be an integer. It is indeed obtained from the continuous-time limit of the ARMA process that describes the average forecast error. But the result, as stated, holds for the true, discrete-time process.

objects we characterized analytically in this section. The latter are the exact counterparts of the empirical objects we documented in Section 3. Clearly, the qualitative pattern is the same regardless of whether we look at one- or three-step ahead forecasts (the same of course applies to other choices of horizon). And the key observation is that *only* with the combination of slow learning and over-extrapolation can the theory generate a sign reversal for the aggregate forecast errors, or average forecasts that undershoot initially and overshoot later on.

In the data, we found that both the impulse response of unemployment forecast errors to the "unemployment shock" and the impulse response of inflation forecast errors to the "inflation shock" displayed the kind of sign reversal seen in the bottom row of Figure 8. Under the lens of the theory, this is strong evidence of both over-extrapolation, or $\hat{\rho} > \rho$, and slow enough learning, or $\hat{\tau}$ small enough.

5.5 Reconciling theory and facts, and an identification strategy

To put everything together, recall the empirical patterns that we found in Section 3:

- For unemployment, $K_{\text{CG}} > 0$, $K_{\text{BGMS}} > 0$, and $K_{\text{IRF}} \in (1, \infty)$.
- For inflation, $K_{\text{CG}} > 0$, $K_{\text{BGMS}} < 0$, and $K_{\text{IRF}} \in (1, \infty)$.

That is, the only essential difference between the two cases is that K_{BGMS} switches signs.

Regardless, the following properties hold by implication of Propositions 1 and 2:

- The last fact, $K_{\text{IRF}} \in (1, \infty)$, necessitates both over-extrapolation, $\hat{\rho} > \rho$, and slow enough learning, or a value for $\hat{\tau}$ not too high.
- Given that some over-extrapolation is needed, the first fact, $K_{\text{CG}} > 0$, puts an additional upper bound on $\hat{\tau}$.
- Suppose the econometrician knows ρ (say, by observing the true outcome z_t and identifying its true persistence). Then, because K_{IRF} and K_{CG} are only functions of $(\hat{\tau}, \rho, \hat{\rho})$ and not of τ, whereas K_{BGMS} is also a function of τ, the econometrician can use K_{IRF} and K_{CG} to identify $(\hat{\tau}, \hat{\rho})$ regardless of the values of K_{BGMS}.
- With $(\hat{\tau}, \hat{\rho})$ identified as above, τ can be chosen so as to match the value of K_{BGMS}.
- This procedure necessarily produces $\hat{\rho} > \rho$, $\hat{\tau} \in (0, \infty)$ and $\tau \in (0, \infty)$ for both cases.
- In the case of unemployment, it also produces $\hat{\tau} < \tau$, or underconfidence, because this is strictly needed in order to generate $K_{\text{BGMS}} > 0$ in the presence of over-extrapolation.
- In the case of inflation, on the other hand, $\hat{\tau}$ could be either lower or higher than τ, depending on how large the over-extrapolation and how negative K_{BGMS}.

This summarizes how the facts fit in our framework, and how the empirical moments we have documented provide identification for the underlying deep belief parameters.

In the next section, we will show how these lessons can be adapted to a GE context. On the one hand, we will explain why the two-step, "triangular" identification strategy proposed above—first get the key parameters $(\hat{\tau}, \hat{\rho})$ and the counterfactual of interest from Facts 1 and 3 alone, then get τ as a "residual" from Fact 2—remains valid in a GE context. On the other hand, we will show how that the identification of $(\hat{\tau}, \hat{\rho})$ and the counterfactual of interest crucially depend, not only on these facts, but also on parameters that govern GE feedbacks, such as the slope of the Keynesian cross. On the way, we also discuss lessons for two types of imperfect expectations that we have not addressed so far, namely Level-k Thinking and cognitive discounting.

6 Into the Wilderness: Imperfect Expectations in GE

We just argued, using a simple framework that abstracted from the endogeneity of inflation and unemployment, that the data require the combination of dispersed private information with a departure from the pillar of rational expectations. The specific required deviation was perceived over-persistence in fundamentals leading to over-extrapolation (and perhaps over- or underconfidence, too). Our ultimate goal is to integrate these insights into the New Keynesian model and quantify their importance. But to do this, we must first understand how imperfect expectations matter in a GE context. This brings us squarely into the "wilderness" of equilibrium fixed-points with noisy and non-rational expectations.

To extract a clear structure out of this wilderness, in this section we accomplish the following tasks. We embed the previously-introduced model of imperfect expectations into a simplified version of the New Keynesian model (with perfectly rigid prices) and analytically characterize the fixed point between expectations and outcomes. We use this to illustrate how the imperfection in expectations influences the dynamic response of the economy to monetary policy or demand shocks. We show in detail how our framework helps capture under the same umbrella the diverse set of theories mentioned in the Introduction. And we finally clarify which moments of the expectations data are most relevant for quantifying the bite of the frictions.

6.1 The Keynesian cross, with and without perfect expectations

When prices are completely rigid, the New Keynesian model boils down to the following equations:

$$y_t = c_t \tag{18}$$

$$c_t = -\varsigma r_t + \mathbb{E}_t^*[c_{t+1}] + \varepsilon_t \tag{19}$$

where y_t is aggregate output, c_t is aggregate spending, r_t is the nominal interest rate (also the real one since prices are completely rigid), ς is the EIS, ε_t is an exogenous discount factor shock (demand shock), and $\mathbb{E}_t^*$ is the *rational* expectation of the representative agent. The first condition is market clearing in the goods market. The second condition is the Euler condition of the representative consumer, or the Dynamic IS curve.

The textbook derivation of this condition is deceptively simple: by imposing a representative consumer, it obscures the GE interaction of multiple consumers. To see this, let us instead allow consumers

to different, and potentially irrational, expectations. Following the same steps as in Angeletos and Lian (2018), one can obtain the following modern version of the Keynesian cross:

$$c_t = \beta \sum_{k=0}^{\infty} \beta^k \overline{\mathbb{E}}_t[-\varsigma r_t + \varepsilon_t] + (1-\beta) \sum_{k=0}^{\infty} \beta^k \overline{\mathbb{E}}_t[y_{t+k}] \tag{20}$$

where $\overline{\mathbb{E}}_t$ denote the average expectation in period t and β is the subjective discount factor. This condition follows directly from aggregating the log-linearized optimal consumption function and aggregating. The second term captures the consumers' present discounted value income, as in the Permanent Income Hypothesis (PIH).

To see more clearly how (20) captures the Keynesian cross, let $Y = \sum_{k=0}^{\infty} \beta^k \overline{\mathbb{E}}_t[y_{t+k}]$ be the average, possibly irrational, expectation of permanent income. We can then read the above condition as $c = a + bY$, where $a \equiv \sum_{k=0}^{\infty} \beta^k \overline{\mathbb{E}}_t[-\varsigma r_t + \varepsilon_t]$ is the intercept of the Keynesian cross and $b \equiv (1-\beta)$ is its slope, or equivalently the marginal propensity to consume out of income (MPC).[26]

For our purposes, it is therefore best to think of β in condition (20) as an inverse measure of the MPC, or the slope of the Keynesian cross. But why does this object drops out from the textbook version of the Dynamic IS curve? Because of "perfect" expectations.

With full information and rational expectations, the average subjective expectation in the population, $\overline{\mathbb{E}}_t$, can be replaced by that of a single, representative, rational agent. One can then apply the Law of Iterated Expectations on condition (20) to reduce it to condition (19). The parameter β, and by the same token the MPC and the slope of the Keynesian cross, then drop out of the picture. But away from that benchmark, one is "stuck" with condition condition (20), and the slope of the Keynesian cross remains a potentially crucial determinate of the aggregate dynamics.

Angeletos and Lian (2018) and Angeletos and Huo (2019) have explored the implications of this insight away from full information. Even if one preserves the pillar of rational expectations, one obtains two distortions in the aggregate spending dynamics—myopia towards the future and anchoring to the past—that both increase not only with the level of noise but also with the slope of the Keynesian cross. In the sequel, we extend their analysis to the richer model of "imperfect expectations" introduced here, explain how the resulting framework helps proxy for additional departures from rational expectations such as cognitive discounting and level-k thinking, and, last but not least, connect to the empirical evidence reported in the first part of the paper.

[26] In the textbook version of the PIH and the New Keynesian model alike, the MPC equals the steady-state interest rate. But consider an OLG version of the New Keynesian model, along the lines of Del Negro, Giannoni, and Patterson (2015) and Piergallini (2007). In each period, a consumer remains alive with probability $\chi \in (0,1]$; with the remaining probability, he dies and gets replaced by a new consumer; and markets are complete, inclusive of annuities. In this case, condition (20) holds with β replaced by $\beta\chi$. By varying χ, we can thus vary the slope of the Keynesian cross, or the strength of the GE feedback, holding constant the "true" discount factor. Furthermore, as shown in **?**, χ can be recast as the probability of binding liquidity constraints. In this sense, a lower value for $\beta\chi$ can be interpreted of as tighter consumer credit.

6.2 Belief structure

The assumptions made about beliefs are essentially the same as before. The only difference is that we now let the imperfection in expectations feed into actual behavior (and vice versa).

Let $\xi_t \equiv -\varsigma r_t + \varepsilon_t$. We henceforth treat ξ_t as an exogenous process and concentrate on how imperfect expectations influence the response of aggregate spending to innovations in ξ_t. In particular, we let ξ_t follow an AR(1) process:

$$\xi_t = \frac{1}{1-\rho\mathbb{L}}\eta_t, \tag{21}$$

where $\rho \in (0,1)$ is a persistence parameter, $\mathbb{L}$ is the lag operator, and η_t is an independently, identically, and normally distributed innovation. A positive η_t can be interpreted as an expansionary momentary policy or an expansionary demand shock.

Consistent with Section 3.1, we let each consumer observe only a noisy private signal of ξ_t, the true precision of which is given by $\tau > 0$. In particular, the signal received in period t is

$$s_{i,t} = \xi_t + \frac{u_{i,t}}{\sqrt{\tau}},$$

where $u_{i,t} \sim N(0,1)$ is idiosyncratic Gaussian noise, and the information of the consumer in period t is given by the history of this signal up to, and including, period t.

As in Section 5, we instead add two behavioral twists. First, we let consumers' subjective perception of the precision of their information be some $\hat{\tau} > 0$, where $\hat{\tau}$ may differ from τ. And second, we let their subjective perception of the persistence of the underlying impulse be some $\hat{\rho} \in (0,1)$, where $\hat{\rho}$ may differ from ρ.

6.3 Solving the fixed point

As already mentioned, Angeletos and Huo (2019) have solved the fixed point of such a beauty contest under the restriction of rational expectations, or $\hat{\rho} = \rho$ and $\hat{\tau} = \tau$. The following proposition extends their solution to the present environment, with the two behavioral twists.

Proposition 3. *1. The equilibrium exists, is unique, and is such that the aggregate spending and income obey the following law of motion:*

$$y_t = \left(1 - \frac{\vartheta}{\hat{\rho}}\right)\left(1 + \frac{\hat{\rho}-\rho}{1-\hat{\rho}}\right)\left(\frac{1}{1-\vartheta\mathbb{L}}\right)y_t^*, \tag{22}$$

where

$$y_t^* \equiv \left(\frac{1}{1-\rho}\right)\left(\frac{1}{1-\rho\mathbb{L}}\right)\eta_t$$

is the frictionless counterpart and ϑ is a scalar contained in $(0,\hat{\rho})$.

2. The average equilibrium forecasts obey the following law of motion:

$$\bar{\mathbb{E}}_t[y_{t+1}] = \left(1 - \frac{\hat{\lambda}}{\hat{\rho}}\right)\frac{1}{1-\vartheta\hat{\lambda}}\,\frac{\hat{\rho}+\vartheta-\hat{\rho}\vartheta(\mathbb{L}+\hat{\lambda})}{(1-\vartheta\mathbb{L})(1-\hat{\lambda}\mathbb{L})}\left(1-\frac{\vartheta}{\hat{\rho}}\right)\left(1+\frac{\hat{\rho}-\rho}{1-\hat{\rho}}\right)y_t^*$$

where ρy_t^* is the frictionless counterpart and $\mathbb{L}$ is a is a scalar contained in $(0,\vartheta)$.

3. *The scalar ϑ is given by the reciprocal of the largest root of the following cubic:*

$$C(z) \equiv 1 - mpc - \left(1 + (1 - mpc)\left(\hat{\rho} + \frac{1}{\hat{\rho}}\right) + \frac{\hat{\tau}}{\hat{\rho}}\right) z + \left(1 - mpc + \hat{\rho} + \frac{1}{\hat{\rho}} + \frac{\hat{\tau}}{\hat{\rho}}\right) z^2 - z^3. \tag{23}$$

It is thus increasing in $\hat{\rho}$, decreasing in $\hat{\tau}$, invariant in ρ and τ, and increasing in mpc. And the scalar $\hat{\lambda}$, which is the same as that in (11), coincides with the value of ϑ corresponding to $mpc = 0$.

In the rest of this Section, we will expand on the economics behind this result and on its usefulness for both theoretical and empirical purposes. For now, we make the following "technical" remarks.

- With perfect expectations, y_t is given by y_t^*, which is merely a rescaling of ξ_t. So in this case, y_t follows an AR(1) process, with root ρ exogenously fixed by the process of ξ_t. With imperfect expectations, instead, y_t follows an AR(2) process, with roots ρ and ϑ. The latter is endogenous, not only to the belief parameters, but also to the slope of the Keynesian cross. This underscores how the GE feedback between expectations and outcomes, which was absent in our earlier analysis, shapes the persistence of y_t.

- In addition to being the source of such endogenous persistence, the belief frictions also influence the scale of the response of y_t to the underlying monetary policy or demand shocks. This scale effect is captured by the term $\left(1 - \frac{\vartheta}{\hat{\rho}}\right)\left(1 + \frac{\hat{\rho} - \rho}{1 - \hat{\rho}}\right)$ in condition (22) and will be later related to whether the economy displays, in effect, a form of myopia or hyperopia.

- The fact that y_t is, endogenously, an AR(2) process complicates the characterization of the forecasts of y_t and of the corresponding moments ($K_{\text{CG}}, K_{\text{BGMS}}, K_{\text{IRF}}$,etc). Fortunately, this does not upset our earlier interpretation of the documented facts: using the characterization of the forecasts in the second part of Proposition 3, we will be able verify that all the *qualitative* properties stated in Section 5 go through.

- There is, however, a twist, which matters *quantitatively*: because the footprint of the various frictions on the joint equilibrium dynamics of the outcome y_t and the expectations thereof depend on the slope of the Keynesian cross, the theoretical counterparts of the empirical moments we looked at before also depend on it. This is evident in Proposition 3 from the property that the root ϑ, which depends not only on the deep belief parameters but also on the MPC, enters the dynamics of both y_t and the average forecasts thereof

A corollary of the last point is that, from the perspective of econometric identification, the moments of the expectations we documented in the data no more offer a direct measurement of the "extent of friction." This anticipates our later quantitative implementation, which will rely on microeconomic estimates of the MPC (and other relevant GE parameters of the full New Keynesian model) both to extract an estimate of the deep belief parameters from the aforementioned moments and to provide an estimate of the causal effect of the various belief frictions on the actual dynamics.

6.4 Empirical footprint

Building on Proposition 3, we can characterize the empirical footprint of imperfect expectations on equilibrium *behavior* in the following terms.

Proposition 4 (Equilibrium Outcomes)**.** *There exist functions Ω_f and Ω_b such that the equilibrium dynamics of the imperfect-expectations economy is the same as that of a perfect-expectations counterpart in which condition (19), the Euler condition of the representative agent, is modified as follows:*

$$c_t = -r_t + \omega_f \mathbb{E}_t^*[c_{t+1}] + \omega_b c_{t-1} \tag{24}$$

where $\omega_f = \Omega_f(\hat{\tau}, \rho, \hat{\rho}, mpc)$, $\omega_b = \Omega_b(\hat{\tau}, \rho, \hat{\rho}; mpc)$, and $\mathbb{E}_t^*$ is the rational, full-information, expectation operator.

This result extends the observational-equivalence result stated in Proposition 3 of Angeletos and Huo (2019) to the forms of irrationality accommodated here. First, it result offers a bridge to simple representative-agent macro models: ω_b resembles habit persistence, ω_f represents a form of myopia (if $\omega_f < 1$) or hyperopia (if $\omega_f > 1$). And second, it lets us understand the variety of beliefs friction we have accommodate so far, as well as cognitive discounting and level-k thinking (more on this momentarily), in terms of different combinations of the coefficients ω_b and ω_f.

The next result shifts focus from the properties of equilibrium *behavior* to the properties of equilibrium *expectations*.

Proposition 5 (Equilibrium Expectations)**.** *There exist functions $\mathcal{K}_{CG}$, $\mathcal{K}_{BGMS}$, $\mathcal{K}_{IRF}$ and F such that the following properties hold:*

1. *The CG regression coefficient is given by $K_{CG} = \mathcal{K}_{CG}(\hat{\tau}, \rho, \hat{\rho}; mpc)$*
2. *The BGMS regression coefficient is given by $K_{BGMS} = \mathcal{K}_{BGMS}(\tau, \hat{\tau}, \rho, \hat{\rho}; mpc)$*
3. *The IRF of the average forecast errors is given by*

$$\left\{ \frac{\partial \overline{Error}_{t+k}}{\partial \eta_t} \right\}_{k \geq 1} = F(\hat{\tau}, \rho, \hat{\rho}; mpc),$$

where $\overline{Error}_t \equiv y_t - \bar{\mathbb{E}}_{t-1}[y_t]$, and the point it first crosses zero from above is given by

$$K_{IRF} = \mathcal{K}_{IRF}(\hat{\tau}, \rho, \hat{\rho}; mpc).$$

This result and a few related results presented in the sequel facilitate the mapping of the theory to the evidence on expectations. As anticipated, the theoretical counterparts of the moments we documented in the empirical part of the paper are herein now shown to depend, not only on the "deep" belief parameters, but also on the MPC, or the slope of the Keynesian cross. Notwithstanding this point, the following property from Section 5 generalizes: the actual level of noise does not enter any of the moments of the *average* forecasts. This is simply because of the law of large numbers: only the *perceived* level of noise matters.

We now proceed to provide a more detailed intuition for how noise and mis-specification work in equilibrium, generate concrete testable predictions, and relate to the facts we documented earlier on. On the way, we explain how our incorporation of under-extrapolation in a GE context helps proxy for level-k thinking and cognitive discounting, and spell out the testable predictions of these theories, too.

6.4.1 Noise and confidence

Let us first isolate the role of noise, with or without overconfidence. This shuts down over- and under-extrapolation (i.e., it sets $\rho = \hat{\rho}$) and recovers, in effect, the scenario studied in Angeletos and Huo (2019).

Proposition 6 (No Over/under-extrapolation)**.** *Suppose* $\hat{\rho} = \rho \in (0,1)$ *and* $(\tau, \hat{\tau}) \in (0,\infty)^2$ *but, potentially,* $\hat{\tau} \neq \tau$*. Then the following statements are true:*

1. [Myopia and anchoring] $\omega_b > 0$ *and* $\omega_f < 1$*. Furthermore, for given belief parameters,* ω_b *increases with mpc and* ω_f *falls with it.*

2. [Predictability of average forecasts] $K_{CG} > 0$.

3. [Predictability of individual forecasts] $sign(K_{BGMS}) = sign(\tau - \hat{\tau})$.

4. [IRF of forecast errors] $K_{IRF} = \infty$*. That is, the IRF of the average forecast errors is uniformly positive.*

Point 1 is essentially the main result from Angeletos and Huo (2019): the introduction of dispersed private information and higher-order uncertainty is akin to the introduction of two distortions, myopia and anchoring, that both increase with the strength of the GE feedback (which, in the present context, is the slope of the Keynesian cross). We refer the reader to this paper for the robustness of this perspective to richer specifications of the information structure (including noisy public signals, endogenous signals, and sticky information); for a thorough discussion of the distinct roles of first- and higher-order uncertainty; and for a translation in terms of "slowing down GE multipliers."

Points 2 and 3 show that this case restricts $K_{\text{CG}} > 0$ and leaves the sign of K_{BGMS} to be determined by the agents' relative confidence. In particular, the scenario on which Angeletos and Huo (2019), Coibion and Gorodnichenko (2012, 2015), Woodford (2003), and the broader literature on informational frictions have focused on, is herein nested with $\hat{\rho} = \rho$ and therefore $K_{\text{CG}} > 0 = K_{\text{BGMS}}$. Adding over/underconfidence allows the theory to accommodate $K_{\text{BGMS}} \neq 0$, but preserves $K_{\text{CG}} > 0$.

Finally, Point 4 shows that, in the absence of over- or under-extrapolation, the theory cannot explain the "sign-switch" in impulse responses which we showed in the data. This extends Corollary 2 to the GE context.

6.4.2 Over- or under-extrapolation

Let us now isolate the role of over- or underconfidence. That is, we now set $\tau = \hat{\tau} = \infty$ and instead let $\rho \neq \hat{\rho}$. In analogy to Proposition 6, we can prove the following result:

Proposition 7 (No Informational Friction)**.** *Suppose $\tau = \hat{\tau} = \infty$ but, potentially, $\hat{\rho} \neq \rho$. Then the following statements are true:*

1. [Myopia or hyperopia, but no anchoring] $\omega_b = 0$ *and* $\omega_f < 1$.

2. [Predictability] $K_{CG} = K_{BGMS}$ *and* $sign(K_{BGMS}) = sign(\rho - \hat{\rho})$.

3. [Uniform dynamic response] $K_{IRF} = \infty$ *if* $\hat{\rho} < \rho$ *and* $K_{IRF} < 1$ *if* $\hat{\rho} > \rho$. *That is, the IRF of the average forecast errors is either uniformly positive or uniformly negative.*

The first point shows that over- or under-extrapolation *by itself* can never create anchoring, for it operates purely by affecting the forward-looking channel. In particular, under our representation, over-extrapolation alone corresponds to pure myopia, and under-extrapolation corresponds to pure hyperopia.

The second points are the GE counterparts of the related points we made in Section 5: over- or under-extrapolation accommodates predictability of forecast errors, but does not allow their dynamic response to switch sign. And the restriction $K_{\text{CG}} = K_{\text{BGMS}}$ encapsulates a lesson that extends to a much larger class of models that allow for misspecified but common beliefs: such models impose that the predictability of forecast errors at individual-level and aggregate-level regressions are the same. It is only the heterogeneity of beliefs, or information, that allows for the predictability in forecast errors to vary with the level of aggregation.[27]

6.4.3 Higher-order doubts, level-k thinking, and cognitive discounting

We now explain how the form of under-extrapolation we have allowed here proxy for the effects of the following three other kinds of departure from rational expectations:

1. The first is a model of pure "higher-order doubts," without either heterogeneous information or learning. Assume that each consumer observes ξ_t with probability 1 but attaches only probability $q \in (0,1)$ that any other consumers also observes ξ_t; with the remaining probability, any other agent is expected to have her belief about ξ_t reset to the prior. Such a model is the main specification in Angeletos and Sastry (2020).[28]

2. The second is a model of level-k thinking. Assume that a consumer of "level 1" perfectly observes ξ_t but assumes all others consumers a default action $c^d_{i,t} = 0$; an agent of level 2 also perfectly observes ξ_t but assumes all other agents play the level-1 action; and this definition recursively extends for any k > 2. Such models have been used to explain the sluggish, and often incomplete, convergence to Nash equilibrium play in laboratory settings (e.g., Nagel, 1995) and, more recently, agents' expectations formation about "unconventional" policy (e.g., Garcıa-Schmidt and Woodford, 2019; Farhi and Werning, 2019; Iovino and Sergeyev, 2017).

[27] This is true asymptotically: in small samples, there can of course be differences,

[28] See also Angeletos and La'O (2009) for an earlier incarnation within the context of the New Keynesian Philips curve, and Izmalkov and Yildiz (2010) for a close cousin in the context of global games.

3. The third is a model of "cognitive discounting," as proposed by Gabaix (2019). Agents observe the fundamental but have misspecified priors about the process of ξ_t and/or that of y_t. In particular, whenever that the actual laws of motion are

$$\xi_t = \rho\xi_{t-1} + \varepsilon_t \qquad \text{and} \qquad y_t = Ry_{t-1} + D\varepsilon_t,$$

for some constants R and D (to be determined as part of the solution), the agents believe that

$$\xi_t = \hat{\rho}\xi_{t-1} + \varepsilon_t \qquad \text{and} \qquad c_t = \hat{R}c_{t-1} + D\varepsilon_t,$$

with $\hat{\rho} \equiv m\rho$ and $\hat{R} \equiv mR$, for some exogenous scalar $m \in (0,1)$ that represents the degrees of "cognitive discounting" applied when the consumers contemplate the future values of both the fundamental and the outcome.

These models have different methodological underpinnings. The first is grounded in the literature on higher-order beliefs ("forecasting the forecasts of others"). The second is motivated by laboratory experiments documenting behavior inconsistent with Nash equilibrium (or REE) and better described by "shallow" recursive reasoning. The third is related to a literature explaining consumer indifference to non-obvious attributes of the economic environment (Gabaix and Laibson, 2006; Gabaix, 2014).

Despite these differences, all these model impose essentially the same distortion in the forecasts of future economic outcomes. In particular, it is easy to show that all three models impose that the average *subjective* expectation and the corresponding *rational* expectation are connected by the following restriction:

$$\overline{\mathbb{E}}_t[y_{t+1}] = d\mathbb{E}^*_t[y_{t+1}],$$

where $d \in (0,1)$ is a scalar that depends on the "deep" parameter $\zeta \in \{q, \mathrm{k}, m\}$ of the respective model. This scalar measures how much consumers underestimate the future response in the behavior of others and, equivalently, the future response of y_t.

In the first model (higher-order doubts), $d < 1$ is the product of underestimating the *knowledge* of others. In the second (level-k thinking), it is the product of underestimating the *rationality* of others. And in the third (cognitive discounting), it is the product of applying a "cognitive discount" when contemplating the future. But in all cases, the essential friction in beliefs is the same and, by direct implication, the observable effect on behavior is also the same.

But now note that the form of under-extrapolation accommodated in our framework plays the same role as well. Indeed, if we shut down the noise, we can show that

$$\overline{\mathbb{E}}_t[y_{t+1}] = \frac{\hat{\rho}}{\rho}\mathbb{E}^*_t[y_{t+1}].$$

It follows that, for any of the aforementioned three models, we can find a value of $\hat{\rho}$ less than ρ such that our model implies the same effective friction in the expectations. The next result verifies that this logic carries over to the entire set of predictions about outcomes, forecasts, and dynamic responses.

Proposition 8. *Consider any of the three models described above, and let* $\zeta \in \{q, k, m\}$ *denote the key parameter for each. For any of these models and any value for the corresponding* ζ*, there exists some* $\hat{\rho}(\rho, \zeta, mpc) < \rho$*,such that*

1. *The outcomes of the original model is observationally equivalent to our own model without noise* $(\tau = \hat{\tau} = \infty)$ *and under-extrapolation* $(\hat{\rho} < \rho)$.

2. *The forecasts of the original model feature* $K_{CG} = K_{BGMS} > 0$ *and* $K_{IRF} = \infty$*, as in Proposition 7 with under-extrapolation.*

To iterate, the basic idea is that all three models have consumers under-estimate the future response of others, which in turn impact behavior in a similar way as an under-estimation of the persistence of ξ_t. The only subtle difference between all these cases is whether the belief mis-specification operates through both PE and GE considerations, or only through GE.

Under-extrapolation and cognitive discounting operate through *both* channels. By contrast, higher-order doubts and level-k thinking operate *only* through GE considerations, for they introduce a mis-specification in how consumers reason about the behavior of others without a mis-specification in how they reason about exogenous impulses. This means that the mapping from the primitive parameters of these two models (the degree of higher-order doubt, q, or the depth of thinking, k) to the value of $\hat{\rho}$ in our isomorphic model is modulated by the MPC, or the relative importance of GE considerations, but does not change the essence.

We conclude that one can think of the case $\hat{\rho} < \rho$ in our framework as a proxy for both level-k thinking and cognitive discounting. This also makes clear that, at least for our purposes, the theories put forward in Gabaix (2019) and Garcıa-Schmidt and Woodford (2019) are the *antithesis* of those postulated in the finance literature on over-extrapolation (Barberis, Shleifer, and Vishny, 1998; Hong and Stein, 1999; Greenwood and Shleifer, 2014; Gennaioli, Ma, and Shleifer, 2015a). And they are at odds with the evidence presented in the first part of our paper.

First, neither of these theories, at least in the versions postulated so far, helps explain why forecasts eventually over-shoot: such overshooting is prima-facie evidence of over-extrapolation. Second, both of these theories restrict $K_{\text{CG}} = K_{\text{BGMS}} > 0$. But even if they were to be augmented with dispersed private information so as to accommodate $K_{\text{CG}} > K_{\text{BGMS}}$, they would still imply $K_{\text{BGMS}} > 0$, which contradicts the data.[29] This is simply because cognitive discounting and level-k thinking are based on premise that *individual* forecasts are misspecified in the particular direction of underestimating the impact of shocks on future outcomes.[30]

The obvious caveat to this conclusion is that it only applies to the particular evidence we have considered here and may not extend to other contexts. Another caveat is that this conclusion is modulated by

[29] To be precise, $K_{\text{BGMS}} > 0$ contradicts the evidence on forecasts of inflation and of various asset prices (Bordalo et al., 2018), but not those of unemployment.

[30] This also corrects a claim made in Gabaix (2019): the relevant test of cognitive discounting is the contradictory evidence regarding *individual* forecast errors provided in Bordalo et al. (2018) and Kohlhas and Broer (2019), not the seemingly favorable evidence about *average* forecast errors provided in Coibion and Gorodnichenko (2015).

parsimony: given the evidence at hand, we cannot reject the hypothesis that agents over-extrapolate the aggregate shocks at the same time that are sallow thinkers with respect to GE, in manner that the total net effect looks as over-extrapolation.[31] But such a fine distinction is irrelevant for our purposes, insofar as the goal is to provide a simple and parsimonious explanation for the facts.

6.5 The bottom line: the "right" theory and its mapping to the data

The combination of our exploration of the survey evidence on expectations and of our unifying prism into the theoretical literature has suggested that the "winner" among the candidate theories is one that combines over-extrapolation with dispersed private information. The basic intuition can be summarized as follows.

To match the fact that K_{CG} is higher than K_{BGMS}, it is necessary that there is enough idiosyncratic noise, or sufficiently dispersed private information. To match the fact that K_{CG} itself is significantly positive, or relatedly that average forecast errors are positive for a few quarters before turning negative, it is necessary that slow learning overwhelms over-extrapolation over short horizons. But as time passes and people accumulate more information, over-extrapolation takes over, producing the observed overshooting in forecasts, or the reversal of sign in the forecast errors.

In Section 5, we formalized these intuitions in a simpler framework that abstracted from the feedback of expectations to actual behavior. The next result extends the argument to the present GE context. It also describes our "identification strategy," or how to extract the deep belief parameters from the documented moments of the macroeconomic forecasts.

Theorem 1. *Consider the following empirical targets:*

$$K_{CG} > 0 \qquad K_{CG} > K_{BGMS} \qquad \textit{and} \qquad K_{IRF} \in (1, \infty),$$

that is, a coefficient in the CG regression that is both positive and "larger" than the BGMS regression coefficient, and overshooting in the IRF of the average forecasts.

1. *These empirical targets can be met in the theory only if all of the following parameter restrictions hold:*

$$\hat{\rho} > \rho, \qquad \hat{\tau} \in (0, b), \qquad \textit{and} \qquad \tau \in (0, \infty),$$

   for some b small enough relative to $\hat{\rho} - \rho$.

2. *For any given* (ρ, mpc)*, the pair* $(\hat{\tau}, \hat{\rho})$ *that matches the empirical targets for* K_{CG} *and* K_{IRF}*, and that also pins down that actual dynamics of* y_t*, is invariant to* τ.

3. *With the pair* $(\hat{\tau}, \hat{\rho})$ *identified as above, the empirical value for* K_{BGMS} *can be matched by choosing appropriately* τ.

[31] Testing this hypothesis would require using simultaneously data on forecasts of unemployment and inflation (the responses of others) and data on forecasts of empirical proxies of exogenous impulses. But we have difficulty imagine the second type of data, except perhaps in an experimental setting.

The first point establishes the claims made above. The second point highlights that $\hat{\tau}$ and $\hat{\rho}$ can be identified from moments of the average forecasts alone, ignoring moments of the individual forecasts. As noted earlier, this is because the joint dynamics of the aggregate outcome and of the average expectations of it are invariant to τ, or the true level of noise. The latter, or equivalently the difference between $\hat{\tau}$ and τ, can then be identified as a "residual" from K_{BGMS}, or equivalently from the gap between K_{CG} and K_{BGMS}.

This echoes the point made earlier. The evidence on the individual forecasts is crucial for understanding the complete landscape of the departure for FIRE, and for conducting certain counterfactuals, such as the following: what would the aggregate dynamics be if one were to switch off over/underconfidence holding the rest of the frictions constant? But such evidence is not essential for understanding the quantitative importance of the overall departure from FIRE: the counterfactual of restoring FIRE can be constructed using merely evidence on aggregate forecasts.

7 The Macroeconomic Effects of Noise and Over-Extrapolation

In the previous section, we simplified the exposition by assuming that prices are infinitely rigid and monetary policy is unresponsive. We now relax these assumptions, allow expectations be imperfect not only for the consumers but also for the firms, connect the theory to the data.

The model's three equations can then be expressed as follows:

$$
\begin{aligned}
c_{i,t} &= \mathbb{E}_{i,t}\left[-\varsigma \sum_{k=0}^{\infty} (1-\text{mpc})^{k+1}(i_{t+k}-\pi_{t+k+1}) + \text{mpc} \sum_{k=0}^{\infty} (1-\text{mpc})^{k} c_{t+k} + \xi_t^d\right] \\
\pi_{i,t} &= \mathbb{E}_{i,t}\left[\theta \sum_{k=0}^{\infty} (\chi\theta)^k \kappa (c_t + \xi_t^s) + (1-\theta)\sum_{k=0}^{\infty} (\chi\theta)^k \pi_{t+k}\right] \\
i_t &= \phi_\pi \pi_t
\end{aligned}
$$

The first equation is the Dynamic IS Curve, modified to allow for informational frictions and mis-specified beliefs, as in the previous section. The second equation is the corresponding modification of the NKPC. The third equation is the rule for monetary policy. ξ_t^d is the demand shock (a discount-factor shock). ξ_t^s is the supply shock (a monopoly-markup, or cost-push, shock).

We close the model by specifying the shock processes and the belief structures in the same way as before. Using the methods of Angeletos and Huo (2019), we then analytically solve for the equilibrium responses of inflation and consumption as functions of two sets of parameters: the behavioral and policy parameters seen in the above equations $(\varsigma, \text{mpc}, \chi, \theta, \kappa, \phi_\pi)$; the actual and perceived persistence of the shocks $(\rho, \hat{\rho})$; and the perceived precision $(\hat{\tau})$. For the reasons already explained, the actual level of noise (or τ) does not enter the determination of either the aggregate outcomes or the average expectations thereof. [32]

[32] Also note that the relevant GE feedbacks are now three: the Keynesian multiplier, which runs inside the DIS curve and was emphasized in the previous section; the dynamic strategic complementary in the firms' price-setting decisions, which running inside the NKPC and is emphasized in Nimark (2008) and Angeletos and Huo (2019); and the feedback between real interest rates, spending, and inflation, which runs across the DIS and the NKPC and is modulated by monetary policy.

To connect the model to the data, we interpret π_t the quarterly rate of inflation and the negative of y_t as the quarterly rate of unemployment. The first choice requires no justification. The second one is based on the logic that, in our model, y_t coincides with the output gap, which in turn is closely related to unemployment both in richer models and in the data. We next fix the model's behavioral and policy parameters to conventional values, as shown in Table X. We finally pick, for each shock, the values of $\rho, \hat{\rho}$ and $\hat{\tau}$ so as to match as well as possible the evidence reported in Section 4.[33] This yields the parameters values seen in Table 2.

Table 4: Exogenously Set Parameters

Parameter	Description	Value
θ	Calvo prob	0.6
κ	Slope of NKPC	0.02
χ	Discount factor	0.99
mpc	MPC	0.3
ς	IES	1.0
ϕ	Monetary policy	1.5

Table 5: Calibrated Parameters

	$\hat{\rho}$	ρ	τ
Demand shock	0.94	0.80	0.38
Supply shock	0.82	0.57	0.15

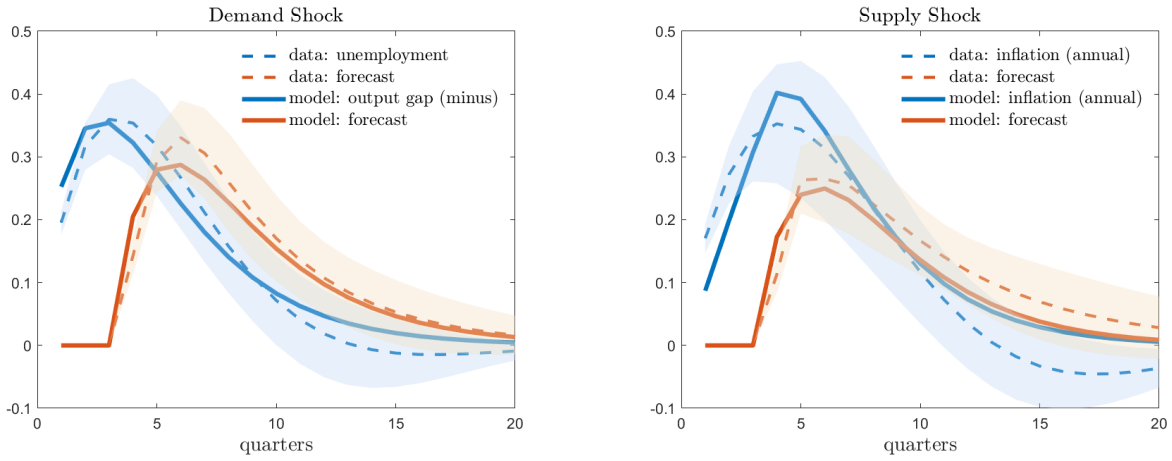
Figure 9 illustrates the model's fit vis-a-vis the empirical IRFs seen earlier in Figure 1. The fit is quite good in the context of the demand shock, but mediocre in the context of the supply shock. This underscores that, although the model has the right *qualitative* ingredients, its *quantitative* performance is not automatic: there is no abundance of degrees of freedom.

We henceforth focus on the demand shock and study two counterfactuals. In the one, we shut down the over-extrapolation, isolating the role of the information friction. In the second, we shut down both frictions, recovering the textbook New Keynesian model. These counterfactuals are illustrated in, respectively, the second and third column of Figure 10. (The first column is the full model, with both frictions, as calibrated above.)

By comparing the second column to the third one, we see that the informational friction alone is the source of both significant dampening and significant persistence relative to the frictionless benchmark. Compared to the textbook model, the informational friction—calibrated to the evidence presented in this paper—decreases the impact of the demand shock on the output gap by about 50% and its impact on

[33] Namely, we minimize the distance between the model-implied IRF of outcomes and forecasts, as well as the term structure of forecasts on impact, from their data counterparts.

Figure 9: Model vs Data



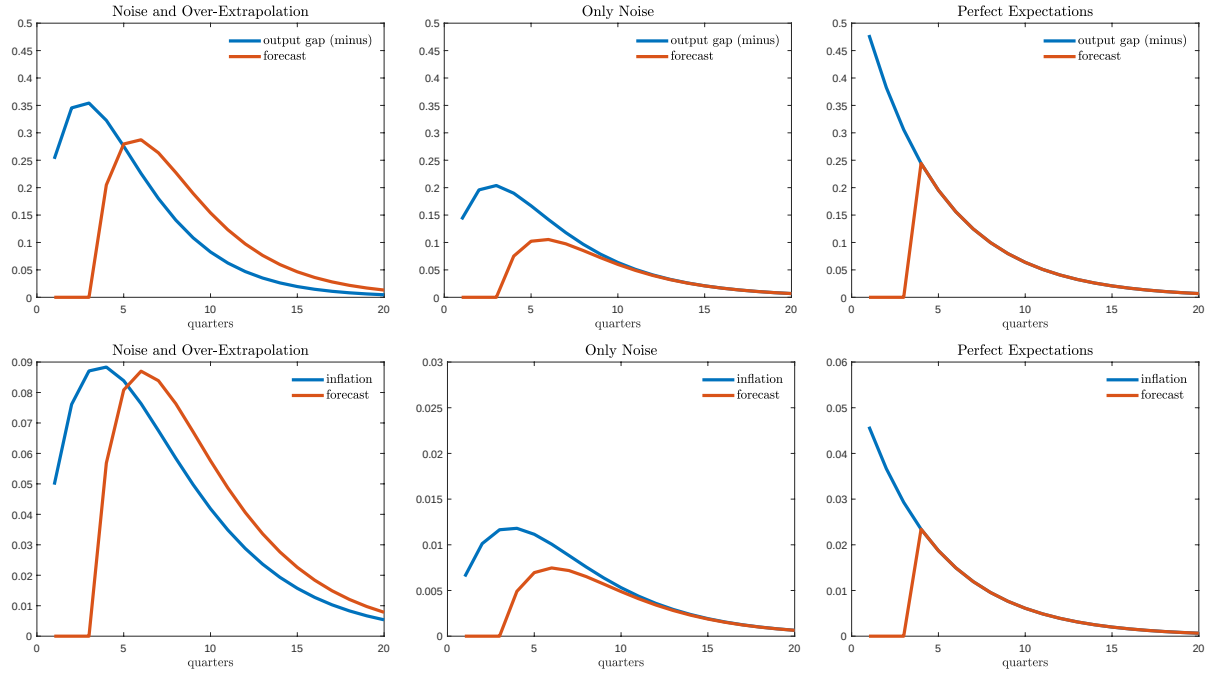
inflation by about 75%. As for the induced persistence, it is quantitatively comparable to that obtained in richer DSGE model with the use of habit persistence in consumption and the hybrid version of the NKPC.

This offers a quantitative assessment, based on forecast evidence, of the common core of Sims (1998, 2003), Mankiw and Reis (2002, 2007), Mackowiak and Wiederholt (2009, 2015), Nimark (2008), Woodford (2003) and a large related literature. But all these works, as well as the related quantitative exercises conducted in Angeletos and Huo (2019), abstract from over-extrapolation. And when over-extrapolation is absent, the model fails to capture our Fact 3: as seen in the second column of Figure 10, the forecasts in the noise-only model do not overshoot. This verifies, once again, that the insight developed in Section 5 about the necessity of combining dispersed information with over-extrapolation extends to the present GE context.

By comparing the first column of Figure 10 to the second one, we then see that the main effect of over-extrapolation on actual outcomes is to amplify their responses to the shock. And while the over-shooting looks "small" in terms of the size of the forecast errors (the gap between the IRFs of forecasts and outcomes after the former have crossed above the latter), its equilibrium footprint is sizable for two reasons. First, a small difference between $\hat{\rho}$ and ρ translates to a large difference in the kind of discounted present values that drive individual behavior. And second, any such belief mistake gets amplified at the aggregate level by GE feedbacks.

So far, we have utilized only evidence on *average* forecasts, ignoring the kind of individual-level evidence that are the focus of Bordalo et al. (2018), Kohlhas and Broer (2019) and Fuhrer (2018). As previously explained, this is because, under the prism of the theories we have considered in this paper, such evidence helps pin down a "residual" parameter that does not enter the dynamics of either the aggregate outcomes or the average forecasts. Such evidence are therefore not strictly needed for the counterfactuals conducted above. But they were useful in pinning down the right model of beliefs. And in the present context, they are essential for conducting *another* counterfactual of interest: what would happen if agents did not suffer from over- or under-confidence. `[to be added]`

Figure 10: Counterfactuals (Demand Shock)



Needless to say, the above quantitative findings should not be taken too seriously. The exercise conducted here is based not only on an overly simplified model but also on imperfect evidence. Ideally, we would like to have evidence about the expectations of firms and consumers, as opposed to those of professional forecasters. Although firms and consumers are almost surely less informed than professional forecasters, it is unclear which class of agents is more prone to over-extrapolation. And while the evidence we presented from the University of Michigan Survey of Consumer Sentiment offers some support for over-extrapolation, the qualitative nature of the questions in this survey preclude the kind of quantitative exercise conducted here.

Last but not least, the most relevant expectations in practice are presumably those that consumers form about their income and their mortgage rates, or those that firms form about the demand for their products and their production and financing costs—but comprehensive, time-series evidence about this kind of expectations is missing.

We therefore view the above exercise only as an illustration of how the marriage of theory with superior evidence on expectations could facilitate a quantitative evaluation of how "imperfect expectations" matters for business cycles.

8 Conclusion

Where are we in the "wilderness" of imperfect expectations? This paper organized theory and empirical evidence from surveys of macroeconomic expectations to show how to best answer that question, taking into account the possibility for multiple competing distortions in expectations formation and the endo-

geneity of outcomes to expectations in general equilibrium.

Our overall conclusion is that, at least to explain the fluctuations most important for business-cycle fluctuations in unemployment and inflation, the data require a combination of (i) informationalfrictions and (ii) a behavioral tendency to over-extrapolate macroeconomic dynamics. Informational frictions help explain, not only why average forecast under-react to innovations on impact, but also why the forecast errors of one agent tend to be forecastable by the forecast errors of *other* agents. Over-extrapolation, on the other hand, helps explain why, following any given shock, average forecasts eventually over-shoot the actual outcome, or equivalently why forecast errors reverse sign after a few quarters.

This kind of over-shooting was the key new fact presented in this paper. Additional, more direct evidence in favor of over-extrapolation was provided by comparing the subjective persistence, as inferred by the term structure of forecasts, to the objective persistence, as inferred from the dynamic response of the actual outcome. Theories that emphasize under-extrapolation, or close cousins thereof such as cognitive discounting and level-K thinking, are at odds with this kind of aggregate-level evidence, as well as with some of the individual-level evidence presented in Bordalo et al. (2018) and Kohlhas and Broer (2019).

At the same time, our analysis has shed new light on which kind of expectations evidence is most relevant for quantifying the overall distortion in expectations. We argued that, at least within the class of economies studied, the overall distortion is best identified by properties of the *aggregate* forecasts and, in particular, those *conditional* on appropriate shocks. These are the most direct counterparts of the objects that matter in the theory.

Unconditional properties of the aggregate forecasts, such as those considered in Coibion and Gorodnichenko (2015), are informative but non-ideal because they confound the adjustment of beliefs to multiple shocks, each of which may be associated with different imperfections in beliefs. And while properties of the individual forecasts, such as those reported in Bordalo et al. (2018) and Kohlhas and Broer (2019), are essential for understanding the complete landscape of expectations, they drop out of the picture when one focuses on the specific counterfactual of what is the overall departure from full-information, rational expectations.

We completed our contribution by an illustration of these points in the context of the New Keynesian model, augmented with informational frictions and over-extrapolation. In our model, similarly to what we found in the data, people are initially slow to catch on to macro changes but eager to ride onto the narrative that these changes are permanent.

A similar co-existence of under-reaction and over-extrapolation is a "classic fact" for many asset prices (De Bondt and Thaler, 1985; Cutler, Poterba, and Summers, 1991; Lakonishok, Shleifer, and Vishny, 1994). Our findings thus represent a step toward unifying our understanding of imperfect expectations in both macroeconomics and finance. This opens up significant possibilities for more sophisticated quantification and interaction with other important macro mechanisms including financial frictions, investment, and housing. In each of these contexts, expectations are paramount; and the possibility of "initial sluggishness combined with eventual overshooting" in expectations could potentially shed new light on boom-and-bust cycles.

References

Acharya, Sushant, Jess Benhabib, and Zhen Huo. 2017. "The Anatomy of Sentiment-Driven Fluctuations." *Yale mimeo* .

Amador, Manuel and Pierre-Olivier Weill. 2010. "Learning from Prices: Public Communication and Welfare." *Journal of Political Economy* 118 (5):866–907.

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. 2018. "Quantifying Confidence." *Econometrica* 86 (5):1689–1726.

———. 2019. "Business Cycle Anatomy." *American Economic Review* forthcoming.

Angeletos, George-Marios and Zhen Huo. 2019. "Myopia and Anchoring." *MIT mimeo* .

Angeletos, George-Marios and Jennifer La'O. 2009. "Incomplete Information, Higher-order Beliefs and Price Inertia." *Journal of Monetary Economics* 56:19–37.

———. 2013. "Sentiments." *Econometrica* 81 (2):739–779.

Angeletos, George-Marios and Chen Lian. 2016. "Incomplete Information in Macroeconomics: Accommodating Frictions in Coordination." *Handbook of Macroeconomics* 2:1065–1240.

———. 2018. "Forward Guidance without Common Knowledge." *American Economic Review* 108 (9):2477–2512.

Angeletos, George-Marios and Karthik A. Sastry. 2020. "Managing Expectations: Instruments vs. Targets." *MIT mimeo* .

Barberis, Nicholas, Andrei Shleifer, and Robert Vishny. 1998. "A model of investor sentiment." *Journal of Financial Economics* 49 (3):307 – 343.

Barsky, Robert B. and Eric R. Sims. 2011. "News shocks and business cycles." *Journal of Monetary Economics* 58 (3):273 – 289.

Benhabib, Jess, Pengfei Wang, and Yi Wen. 2015. "Sentiments and Aggregate Demand Fluctuations." *Econometrica* 83 (2):549–585.

Bergemann, Dirk and Stephen Morris. 2013. "Robust Predictions in Games with Incomplete Information." *Econometrica* 81 (4):1251–1308.

Bhandari, Anmol, Jaroslav Borovička, and Paul Ho. 2019. "Survey Data and Subjective Beliefs in Business Cycle Models." University of Minnesota, New York University and Federal Reserve Bank of Richmond.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. 2018. "Over-reaction in Macroeconomic Expectations." *NBER Working Paper No. 24932* .

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2017. "Diagnostic Expectations and Credit Cycles." *The Journal of Finance* 73 (1):199–227.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo. 2016. "Understanding Booms and Busts in Housing Markets." *Journal of Political Economy* 124 (4):1088–1147.

Caballero, Ricardo and Alp Simsek. 2017. "A Risk-centric Model of Demand Recessions and Macroprudential Policy." miméo, MIT.

Carroll, Christopher D. 2001. "The Epidemiology of Macroeconomic Expectations." Working Paper 8695, National Bureau of Economic Research.

Chahrour, Ryan and Gaetano Gaballo. 2018. "Learning from Prices: Amplification and Business Fluctuations." *Boston College/ECB mimeo* .

Christiano, Lawrence J., Roberto Motto, and Massimo Rostagno. 2014. "Risk Shocks." *American Economic Review* 104 (1):27–65.

Coibion, Olivier and Yuriy Gorodnichenko. 2012. "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy* 120 (1):116–159.

———. 2015. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review* 105 (8):2644–78.

Coibion, Olivier, Yuriy Gorodnichenko, and Saten Kumar. 2018. "How Do Firms Form Their Expectations? New Survey Evidence." *American Economic Review* 108 (9):2671–2713.

Coibion, Olivier, Yuriy Gorodnichenko, and Tiziano Ropele. 2019. "Inflation Expectations and Firm Decisions: New Causal Evidence*." *The Quarterly Journal of Economics* 135 (1):165–219.

Cutler, David M., James M. Poterba, and Lawrence H. Summers. 1991. "Speculative Dynamics." *The Review of Economic Studies* 58 (3):529–546.

D'Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber. 2019. "Exposure to Daily Price Changes and Inflation Expectations." Working Paper 26237, National Bureau of Economic Research.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam. 1998. "Investor Psychology and Security Market Under- and Overreactions." *The Journal of Finance* 53 (6):1839–1885.

Das, Sreyoshi, Camelia M Kuhnen, and Stefan Nagel. 2020. "Socioeconomic status and macroeconomic expectations." *The Review of Financial Studies* 33 (1):395–432.

De Bondt, Werner FM and Richard Thaler. 1985. "Does the stock market overreact?" *The Journal of Finance* 40 (3):793–805.

Del Negro, Marco, Marc P Giannoni, and Christina Patterson. 2015. "The Forward Guidance Puzzle." *FRB of New York mimeo* .

Eusepi, Stefano and Bruce Preston. 2011. "Expectations, Learning, and Business Cycle Fluctuations." *American Economic Review* 101 (6):2844–72.

Evans, George W and Seppo Honkapohja. 2012. *Learning and expectations in macroeconomics*. Princeton University Press.

Farhi, Emmanuel and Iván Werning. 2019. "Monetary Policy, Bounded Rationality, and Incomplete Markets." *American Economic Review* 109 (11):3887–3928.

Fuhrer, Jeffrey C. 2018. "Intrinsic expectations persistence: evidence from professional and household survey expectations." Working Papers 18-9, Federal Reserve Bank of Boston.

Fuster, Andreas, David Laibson, and Brock Mendel. 2010. "Natural Expectations and Macroeconomic Fluctuations." *Journal of Economic Perspectives* 24 (4):67–84.

Gabaix, Xavier. 2014. "A Sparsity-Based Model of Bounded Rationality." *Quarterly Journal of Economics* 129 (4):1661–1710.

———. 2019. "A Behavioral New Keynesian Model." *Harvard mimeo* .

Gabaix, Xavier and David Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets*." *The Quarterly Journal of Economics* 121 (2):505–540.

Galí, Jordi. 1999. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review* 89 (1):249–271.

Garcıa-Schmidt, Mariana and Michael Woodford. 2019. "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis." *American Economic Review* 109 (1):86–120.

Geanakoplos, John. 2010. "The leverage cycle." In *NBER Macroeconomics Annual 2009, Volume 24*. University of Chicago Press, 1–65.

Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer. 2015a. "Expectations and investment." In *NBER Macroeconomics Annual 2015, Volume 30*. University of Chicago Press.

———. 2015b. "Expectations and Investment." *NBER working paper 21260* .

Greenwood, Robin and Andrei Shleifer. 2014. "Expectations of returns and expected returns." *Review of Financial Studies* 27 (3):714–746.

Guo, Hongye and Jessica A Wachter. 2019. "'Superstitious' Investors." Working Paper 25603, National Bureau of Economic Research.

Hamilton, James D. 1996. "This is what happened to the oil price-macroeconomy relationship." *Journal of Monetary Economics* 38 (2):215–220.

Hassan, Tarek A and Thomas M Mertens. 2014. "Information Aggregation in a Dynamic Stochastic General Equilibrium Model." *NBER Macroeconomics Annual* 29 (1):159–207.

Hong, Harrison and Jeremy C. Stein. 1999. "A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets." *The Journal of Finance* 54 (6):2143–2184.

Ilut, Cosmin L and Martin Schneider. 2014. "Ambiguous Business Cycles." *American Economic Review* 104 (8):2368–2399.

Iovino, Luigi and Dmitriy Sergeyev. 2017. "Quantitative Easing without Rational Expectations." *Work in progress* .

Izmalkov, Sergei and Muhamet Yildiz. 2010. "Investor sentiments." *American Economic Journal: Microeconomics* 2 (1):21–38.

Jordà, Òscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95 (1):161–182.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti. 2010. "Investment Shocks and Business Cycles." *Journal of Monetary Economics* 57 (2):132–145.

Kohlhas, Alexandre and Tobias Broer. 2019. "Forecaster (Mis-)Behavior." *IIES miméo* .

Lakonishok, Josef, Andrei Shleifer, and Robert W. Vishny. 1994. "Contrarian Investment, Extrapolation, and Risk." *The Journal of Finance* 49 (5):1541–1578.

Lorenzoni, G. 2010. "Optimal Monetary Policy with Uncertain Fundamentals and Dispersed Information." *The Review of Economic Studies* 77 (1):305–338.

Lorenzoni, Guido. 2009. "A Theory of Demand Shocks." *American Economic Review* 99 (5):2050–84.

Lucas, Robert E. Jr. 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory* 4 (2):103–124.

Mackowiak, Bartosz and Mirko Wiederholt. 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review* 99 (3):769–803.

———. 2015. "Business Cycle Dynamics under Rational Inattention." *Review of Economic Studies* 82 (4):1502–1532.

Malmendier, Ulrike and Stefan Nagel. 2016. "Learning from inflation experiences." *The Quarterly Journal of Economics* 131 (1):53–87.

Mankiw, N Gregory and Ricardo Reis. 2002. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* :1295–1328.

———. 2007. "Sticky information in general equilibrium." *Journal of the European Economic Association* 5 (2-3):603–613.

———. 2011. "Imperfect Information and Aggregate Supply." *Handbook of Monetary Economics* .

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers. 2004. "Disagreement about Inflation Expectations." In *NBER Macroeconomics Annual 2003, Volume 18*, NBER Chapters. National Bureau of Economic Research, Inc, 209–270.

Moore, Don A and Paul J Healy. 2008. "The trouble with overconfidence." *Psychological review* 115 (2):502.

Morris, Stephen and Hyun Song Shin. 2002. "Social Value of Public Information." *American Economic Review* 92 (5):1521–1534.

Nagel, Rosemarie. 1995. "Unraveling in Guessing Games: An Experimental Study." *The American Economic Review* 85 (5):1313–1326.

Nimark, Kristoffer. 2008. "Dynamic Pricing and Imperfect Common Knowledge." *Journal of Monetary Economics* 55 (2):365–382.

Piergallini, Alessandro. 2007. "Real Balance Effects and Monetary Policy." *Economic Inquiry* 44 (3):497–511.

Romer, Christina D. and David H. Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review* 94 (4):1055–1084.

Sargent, Thomas J. 2001. *The conquest of American inflation*. Princeton University Press.

Sims, Christopher A. 1998. "Stickiness." *Carnegie-Rochester Conference Series on Public Policy* 49 (1):317–356.

Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3):665–690.

———. 2010. "Rational Inattention and Monetary Economics." *Handbook of Monetary Economics* 3:155–181.

Sims, Christopher A. and Tao Zha. 2006. "Were There Regime Switches in U.S. Monetary Policy?" *American Economic Review* 96 (1):54–81.

Simsek, Alp. 2013. "Belief disagreements and collateral constraints." *Econometrica* 81 (1):1–53.

Uhlig, Harald. 2003. "What Moves Real GNP?" mimeo.

Veldkamp, Laura L. 2006. "Media frenzies in markets for financial information." *The American Economic Review* :577–601.

———. 2011. *Information choice in macroeconomics and finance*. Princeton University Press.

Woodford, Michael. 2003. "Imperfect Common Knowledge and the Effects of Monetary Policy." *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps* .

———. 2009. "Information-Constrained State-Dependent Pricing." *Journal of Monetary Economics* 56:100–124.

Appendices

A Extra Tables and Figures

Table A.1: Predicting Aggregate Forecast Errors (All Horizons)

$$\text{Error}_{t,k} = K_{CG} \cdot \text{Revision}_{t,k} + \alpha + u_{t,k}$$

	(1)	(2)	(3)	(4)	(5)	(6)
	Unemployment			Inflation		
horizon k	1	2	3	1	2	3
$\text{Revision}_{t,k}$ (K_{CG})	0.384	0.606	0.741	0.649	1.048	1.528
	(0.128)	(0.178)	(0.232)	(0.290)	(0.337)	(0.418)
R^2	0.111	0.143	0.111	0.122	0.200	0.278
Observations	196	196	191	195	195	190

Notes: The dataset is the Survey of Professional Forecasters and the observation is a quarter between Q4-1968 and Q4-2017. All regressions include a constant. Standard errors are HAC-robust, with a Bartlett ("hat") kernel and lag length equal to 4 quarters. The data used for outcomes are first-release ("vintage").

Table A.2: Predicting Aggregate Forecast Errors (1984 to Present)

$$\text{Error}_{t,k} = K_{CG} \cdot \text{Revision}_{t,k} + \alpha + u_{t,k}$$

	(1)	(2)	(3)	(4)	(5)	(6)
	Unemployment			Inflation		
horizon k	1	2	3	1	2	3
$\text{Revision}_{t,k}$ (K_{CG})	0.385	0.657	0.809	-0.100	0.160	0.292
	(0.203)	(0.255)	(0.305)	(0.159)	(0.174)	(0.191)
R^2	0.116	0.195	0.159	0.002	0.005	0.016
Observations	136	136	136	135	135	135

Notes: The dataset is the Survey of Professional Forecasters and the observation is a quarter between Q1-1984 and Q4-2017. All regressions include a constant. Standard errors are HAC-robust, with a Bartlett ("hat") kernel and lag length equal to 4 quarters. The data used for outcomes are first-release ("vintage").

Table A.3: Predicting Aggregate Forecast Errors (Final Release Data)

$$\text{Error}_{t,k} = K_{CG} \cdot \text{Revision}_{t,k} + \alpha + u_{t,k}$$

	(1)	(2)	(3)	(4)	(5)	(6)
	Unemployment			Inflation		
horizon k	1	2	3	1	2	3
$\text{Revision}_{t,k}$ (K_{CG})	0.411	0.612	0.731	0.578	0.991	1.403
	(0.127)	(0.180)	(0.233)	(0.215)	(0.261)	(0.334)
R^2	0.135	0.147	0.108	0.104	0.200	0.249
Observations	199	198	192	199	198	192

Notes: The dataset is the Survey of Professional Forecasters and the observation is a quarter between Q4-1968 and Q4-2017. All regressions include a constant. Standard errors are HAC-robust, with a Bartlett ("hat") kernel and lag length equal to 4 quarters. The data used for outcomes are modern ("final release")

Table A.4: Predicting Aggregate Forecast Errors (Blue Chip Data)

$$\text{Error}_{t,k} = K_{CG} \cdot \text{Revision}_{t,k} + \alpha + u_{t,k}$$

	(1)	(2)	(3)	(4)	(5)	(6)
	Unemployment			Inflation		
horizon k	1	2	3	1	2	3
$\text{Revision}_{t,k}$ (K_{CG})	0.310	0.544	0.804	0.024	0.378	0.618
	(0.129)	(0.213)	(0.231)	(0.204)	(0.188)	(0.205)
R^2	0.091	0.132	0.149	0.000	0.033	0.067
Observations	151	151	150	150	150	149

Notes: The observation is a quarter between Q1-1980 and Q4-2017. Blue Chip publications are matched to quarters by taking the last survey within the quarter (e.g., March for Q1). All regressions include a constant. Standard errors are HAC-robust, with a Bartlett ("hat") kernel and lag length equal to 4 quarters. The data used for outcomes are first-release ("vintage").

Table A.5: Predicting Individual Forecast Errors (All Horizons)

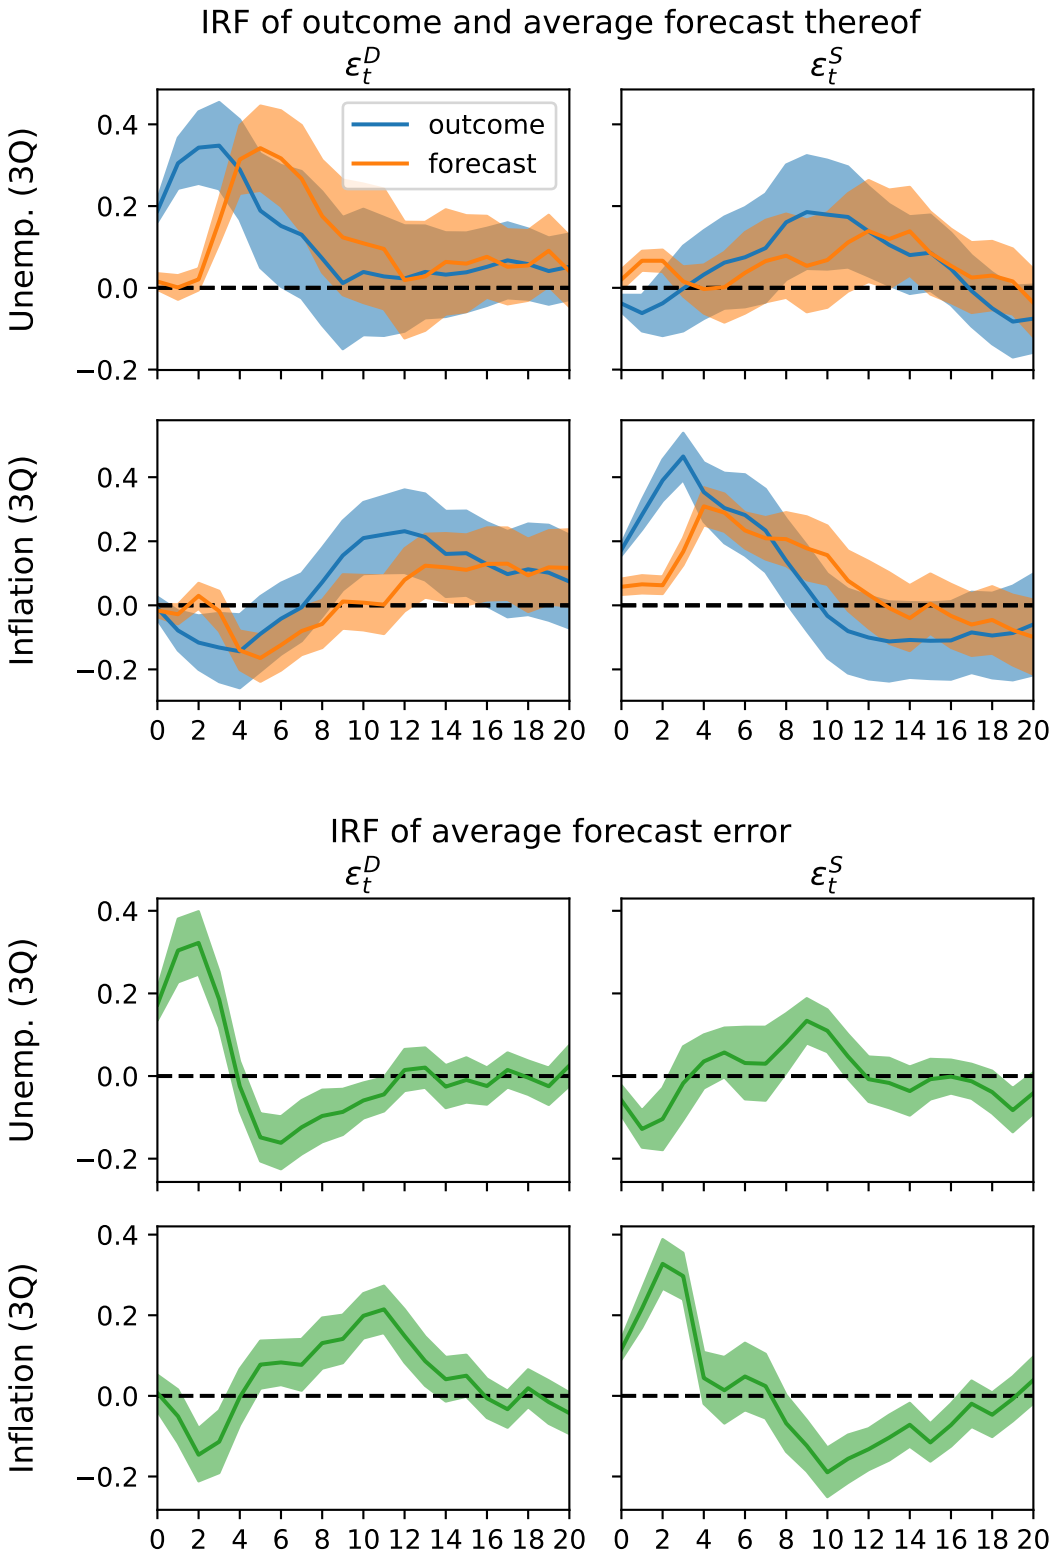
$$\text{Error}_{i,t,k} = K_{BGMS} \cdot \text{Revision}_{i,t,k} + \alpha + u_{i,t,k}$$

	(1)	(2)	(3)	(4)	(5)	(6)
	Unemployment					
horizon k	1		2		3	
Period	1968-2017	1984-2017	1968-2017	1984-2017	1968-2017	1984-2017
$\text{Revision}_{i,t,k}$ (K_{BGMS})	0.186	0.182	0.300	0.322	0.321	0.398
	(0.077)	(0.115)	(0.094)	(0.141)	(0.107)	(0.149)
R^2	0.029	0.030	0.042	0.052	0.028	0.052
Observations	5808	3986	5699	3918	5383	3769

	(7)	(8)	(9)	(10)	(11)	(12)
	Inflation					
horizon k	1		2		3	
Period	1968-2017	1984-2017	1968-2017	1984-2017	1968-2017	1984-2017
$\text{Revision}_{i,t,k}$ (K_{BGMS})	-0.100	-0.439	0.024	-0.360	0.143	-0.263
	(0.084)	(0.045)	(0.098)	(0.044)	(0.123)	(0.054)
R^2	0.004	0.089	0.000	0.053	0.005	0.025
Observations	5496	3779	5458	3745	5147	3643

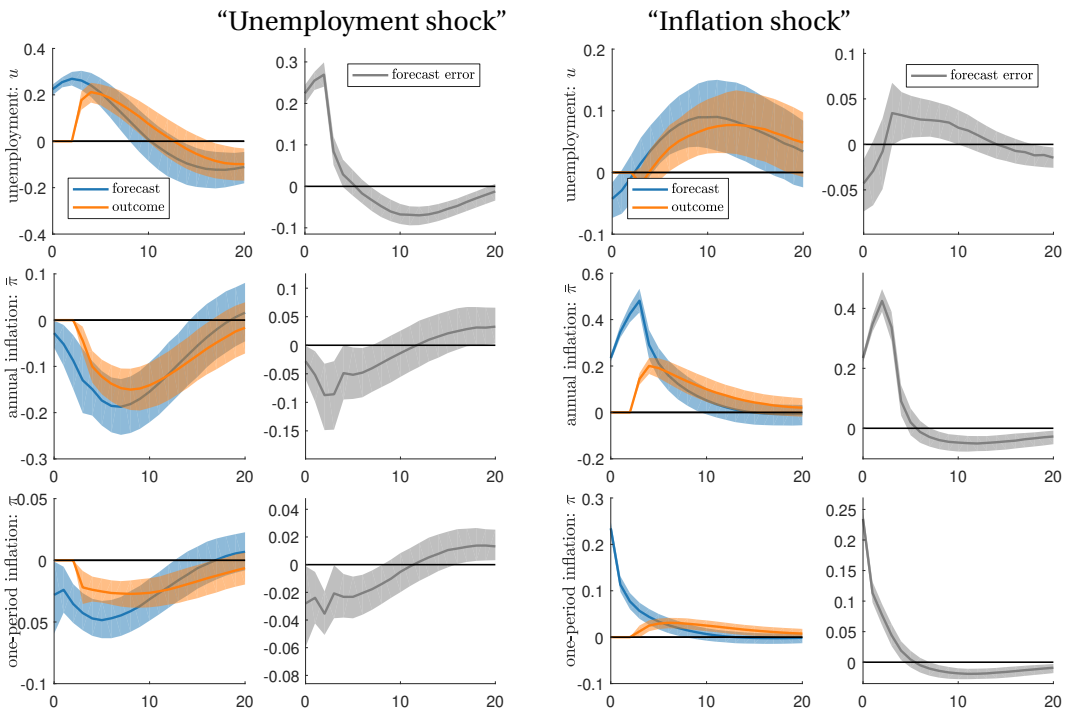
Notes: The observation is a forecaster by quarter between Q4-1968 and Q4-2017. Standard errors are clustered two-way by forecaster ID and time period. Both errors and revisions are winsorized over the sample to restrict to 4 times the inter-quartile range away from the median. The data used for outcomes are first-release ("vintage").

Figure A.1: The Dynamic Response of Unemployment (above) and Inflation (below), All Shock Combinations



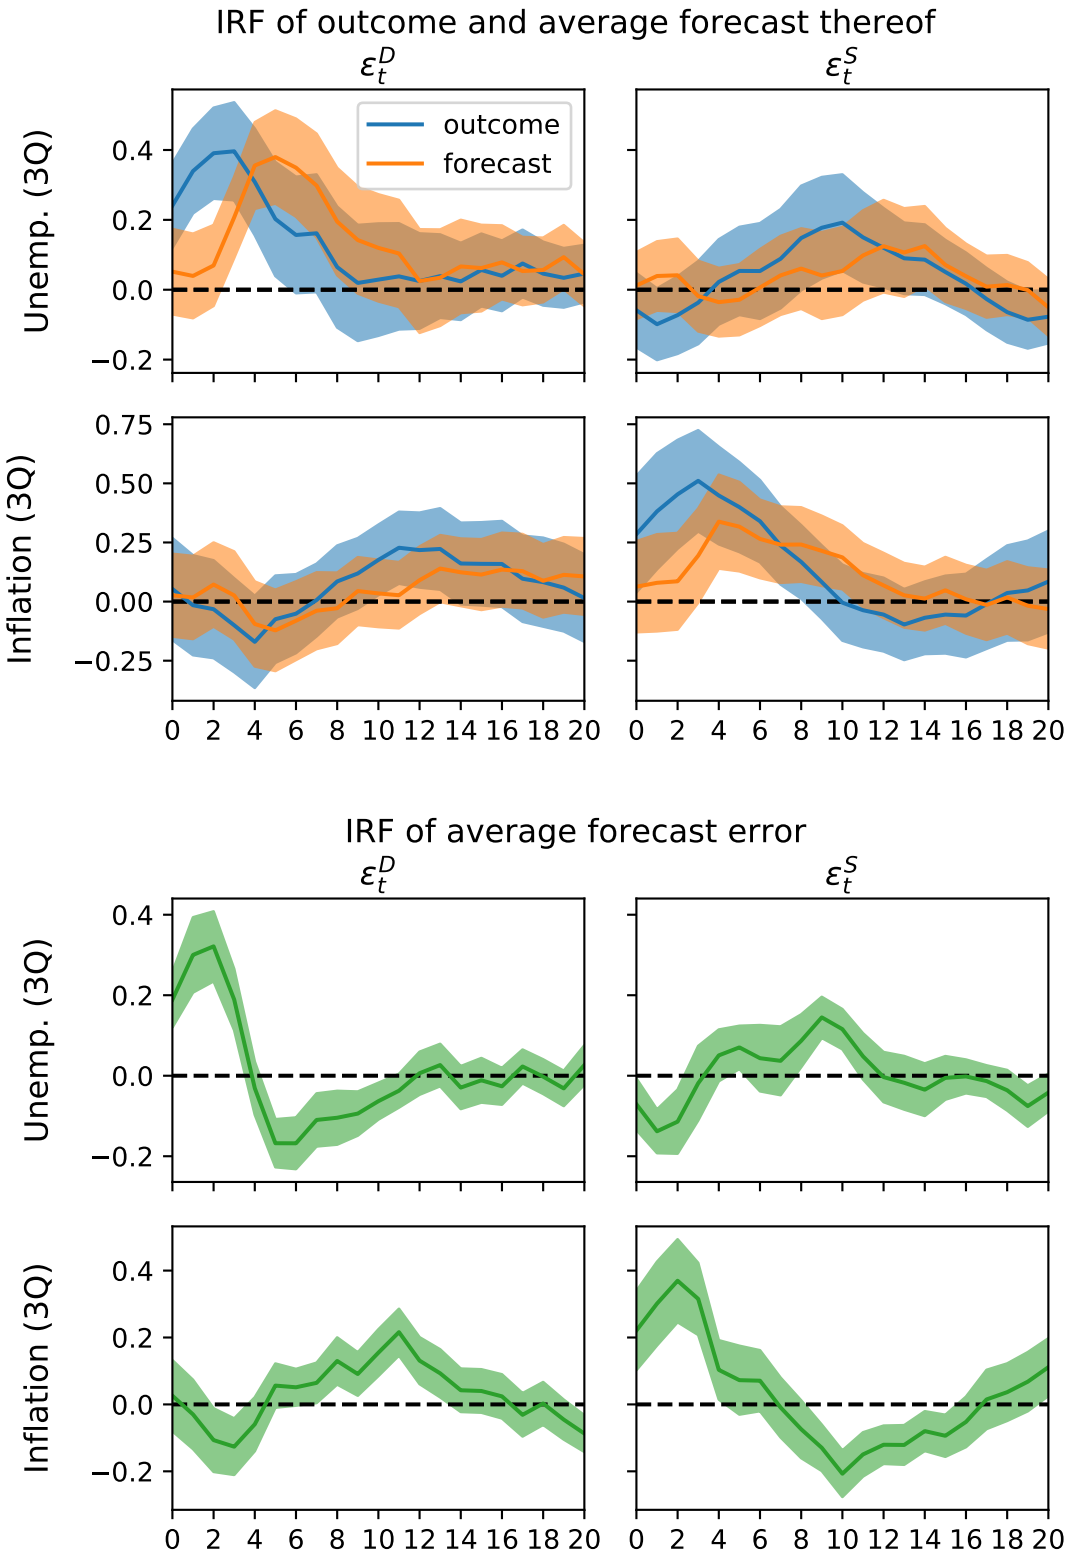
Notes: The sample period is Q4 1968 to Q4 2017. The shaded areas are 68% confidence intervals based on HAC standard errors with a Bartlett ("tent") kernel and 4 lags. In the first row of each panel the outcome is u_t and the forecast is $\bar{\mathbb{E}}_{t-3}[u_t]$; in the second row the outcome is $\pi_{t,t-4}$, or annual inflation, and the forecast is $\bar{\mathbb{E}}_{t-3}[\pi_{t,t-4}]$.

Figure A.2: Dynamic Responses in the Angeletos, Collard, and Dellas (2019) SVAR, All Responses



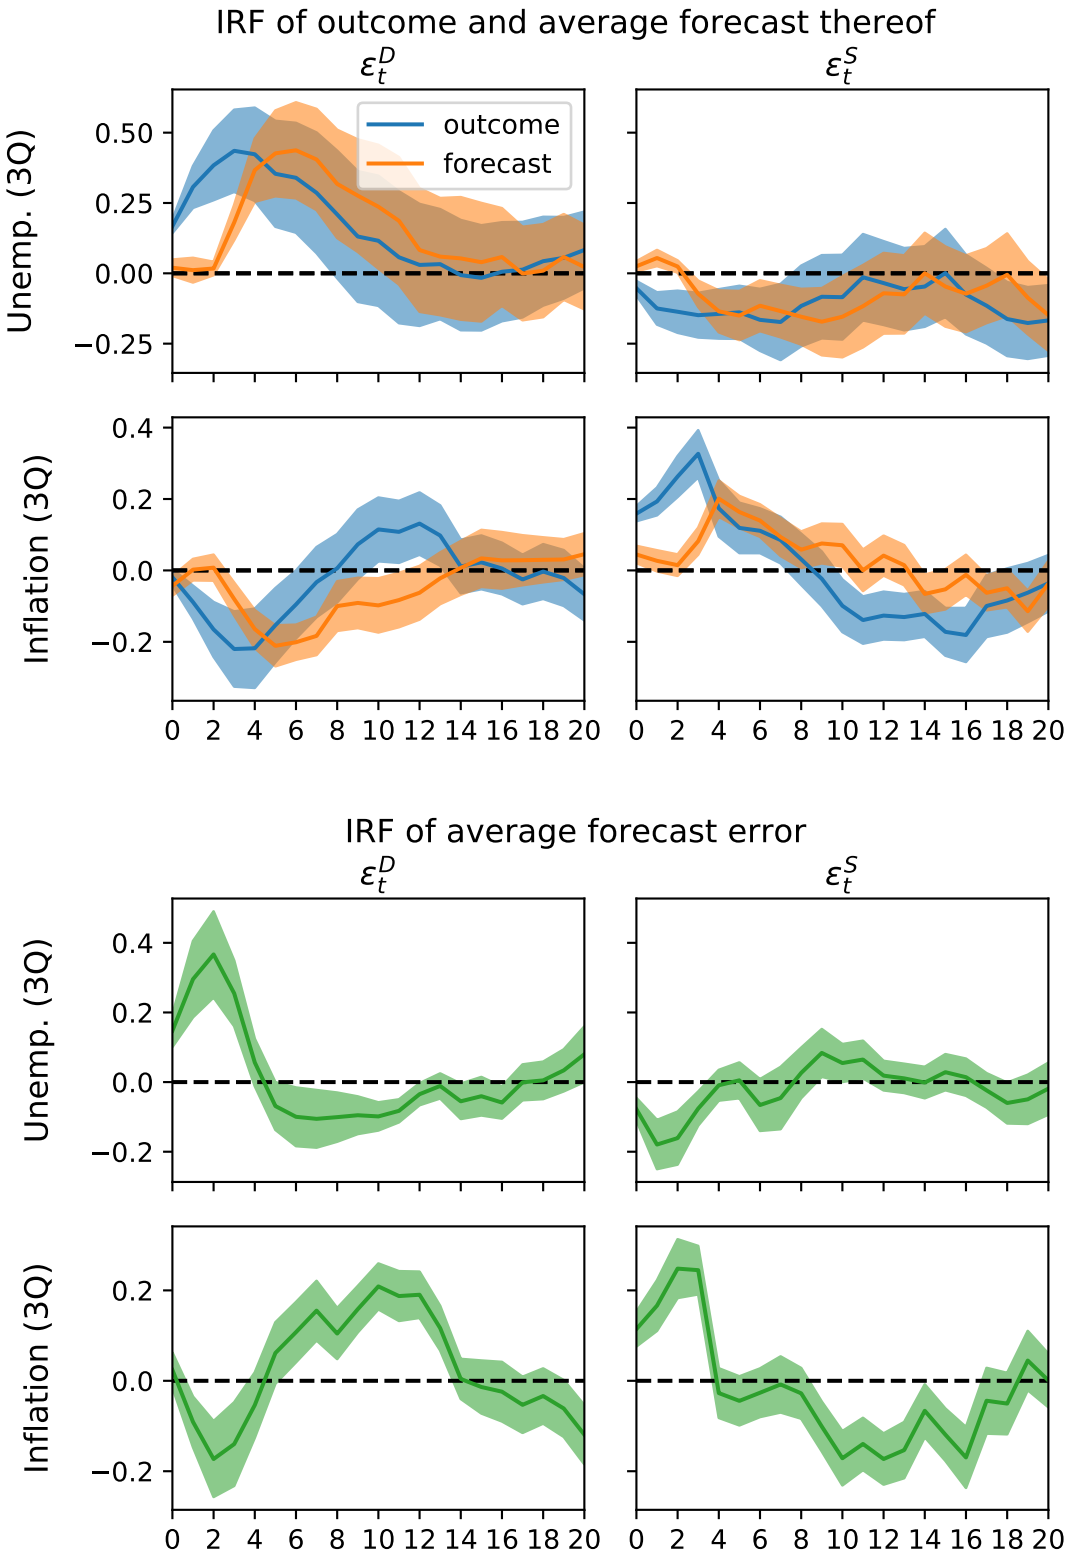
Notes: The sample period is Q4 1968 to Q4 2017. The x-axis denotes quarters from the shock (starting at 0). The shaded areas are 68% high-posterior-density regions and the point estimate is the posterior median. In the first row the outcome is u_t and the forecast is $\bar{\mathbb{E}}_{t-3}[u_t]$; in the second row the outcome is $\pi_{t,t-4}$, or annual inflation, and the forecast is $\bar{\mathbb{E}}_{t-3}[\pi_{t,t-4}]$; and in the last row, the outcome is $\pi_{t,t-1}$, or one-quarter inflation, and the forecast is $\bar{\mathbb{E}}_{t-3}[\pi_{t,t-1}]$. The first column shows the response to a shock that maximizes the business-cycle variation in unemployment; the second for a shock that maximizes the business-cycle variation in GDP deflator inflation.

Figure A.3: The Dynamic Response of Unemployment (above) and Inflation (below), Vintage Data



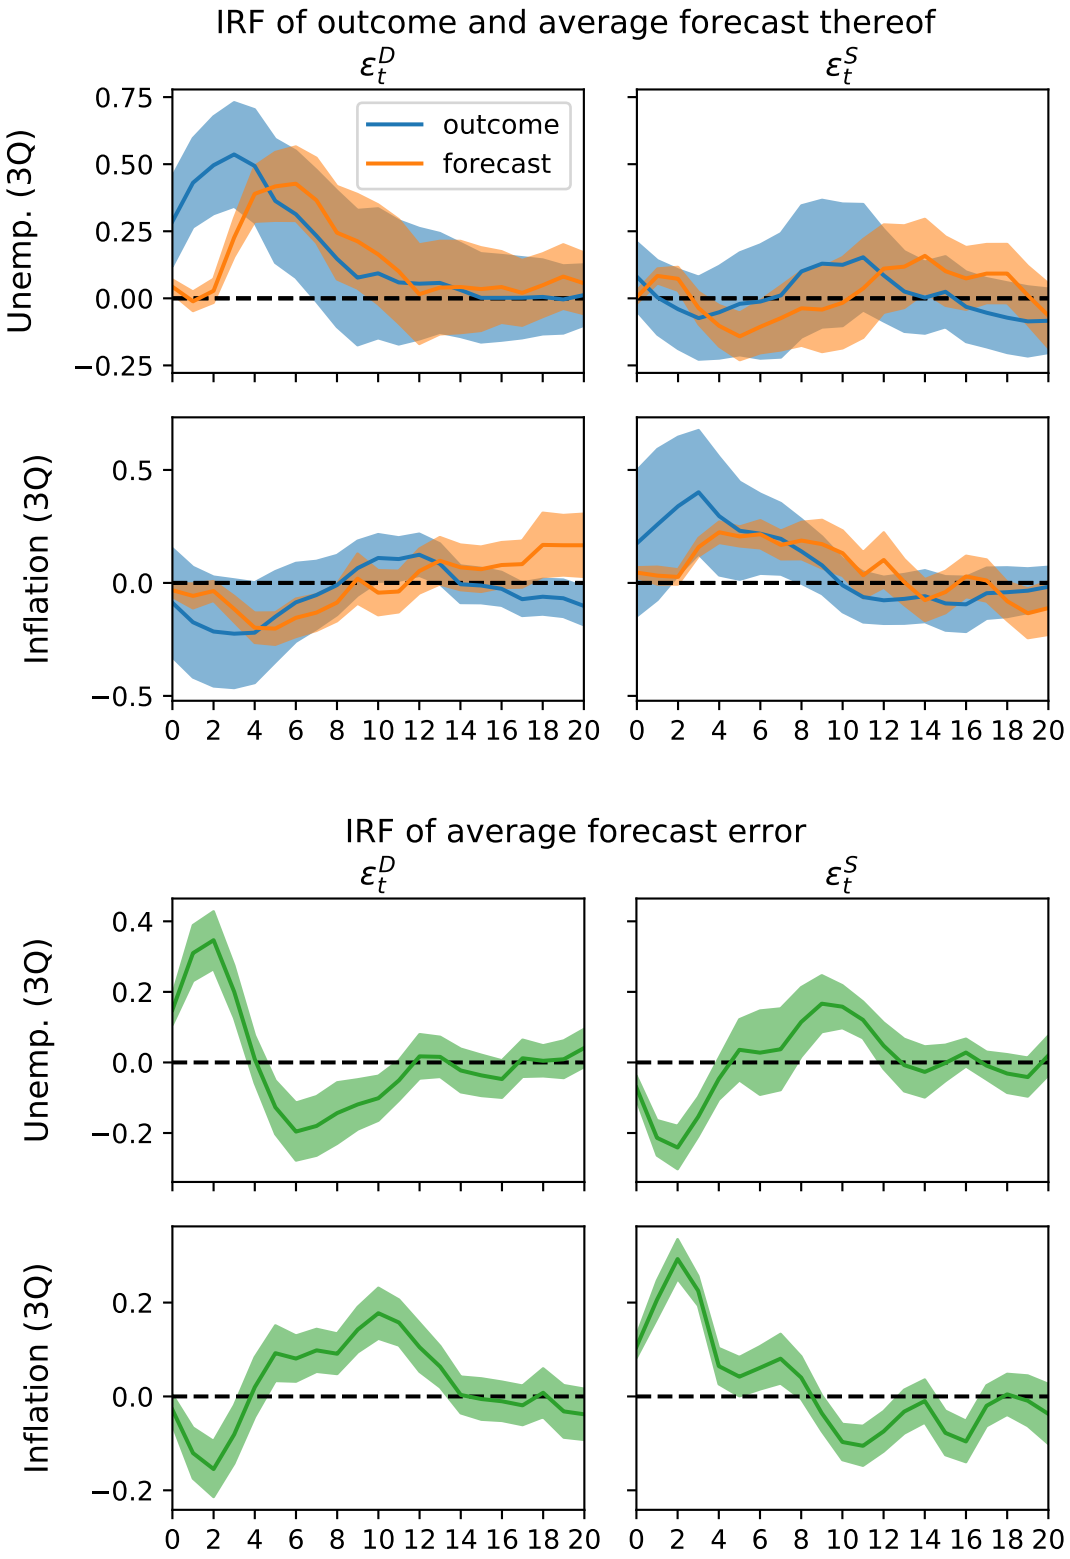
Notes: The sample period is Q4 1968 to Q4 2017. The shaded areas are 68% confidence intervals based on HAC standard errors with a Bartlett ("tent") kernel and 4 lags. In the first row of each panel the outcome is u_t and the forecast is $\bar{\mathbb{E}}_{t-3}[u_t]$; in the second row the outcome is $\pi_{t,t-4}$, or annual inflation, and the forecast is $\bar{\mathbb{E}}_{t-3}[\pi_{t,t-4}]$.

Figure A.4: The Dynamic Response of Unemployment (above) and Inflation (below), Post 1984



Notes: The sample period is Q1 1984 to Q4 2017. The shaded areas are 68% confidence intervals based on HAC standard errors with a Bartlett ("tent") kernel and 4 lags. In the first row of each panel the outcome is u_t and the forecast is $\bar{\mathbb{E}}_{t-3}[u_t]$; in the second row the outcome is $\pi_{t,t-4}$, or annual inflation, and the forecast is $\bar{\mathbb{E}}_{t-3}[\pi_{t,t-4}]$.

Figure A.5: The Dynamic Response of Unemployment (above) and Inflation (below), Blue Chip Data



Notes: The sample period is Q1 1980 to Q4 2017. The shaded areas are 68% confidence intervals based on HAC standard errors with a Bartlett ("tent") kernel and 4 lags. In the first row of each panel the outcome is u_t and the forecast is $\bar{\mathbb{E}}_{t-3}[u_t]$; in the second row the outcome is $\pi_{t,t-4}$, or annual inflation, and the forecast is $\bar{\mathbb{E}}_{t-3}[\pi_{t,t-4}]$.

B Conditional vs. Unconditional Dynamics

Coibion and Gorodnichenko (2012) test models of expectations inertia by estimating the dynamic response of outcomes, forecasts, and forecast errors to shocks, just like this paper does in Section 4. But, while this paper and Coibion and Gorodnichenko (2012) agree about the initial under-reaction of professional forecasters to economic shocks, only the present paper finds robust evidence of the "over-shooting" that we characterize as *Fact 3*. What explains the differences in results, given that our analyses study similar data over a similar time period?

In this section, we will show via a replication of an illustrative main result in Coibion and Gorodnichenko (2012), the response of inflation forecasts and forecast errors to an identified technology shock, that a major difference is estimation methodology—simpler functional forms for the impulse response, while they are parsimonious fits to the data, may not capture all the interesting dynamics (or, for that matter, allow for a very good mapping to what a more structured theory like ours demands).

B.1 Measurement and Methods

We will focus on one main result from Coibion and Gorodnichenko (2012): that inflation expectations respond sluggishly to an inflationary negative supply shock. In this Appendix, we will recreate this fact using the data directly provided by Coibion and Gorodnichenko (2012) for the strongest comparability, although these data are of course essentially identical to those used in our own main analysis.[34] The sample period runs from Q4 of 1974 to Q4 of 2007.

To identify a technology shock, the authors run a four-lag, three-variable VAR with labor productivity, the change in labor hours, and the (one-quarter-ahead) GDP deflator inflation and apply the long-run restrictions introduced by Galí (1999). The estimation period for this VAR covers Q2 1952 to Q3 2007. Finally, to make the shock inflationary, we take the *negative* shock which corresponds to a technological contraction.

Romer and Romer (2004) Impulse Response Estimation. To estimate impulse responses, Coibion and Gorodnichenko (2012) apply the following method due to Romer and Romer (2004). For a given variable x_t (e.g., forecast errors), they estimate the empirical ARMA process via Ordinary Least Squares (OLS):

$$x_t = \alpha + \sum_{p=1}^{P} \gamma_p \cdot x_{t-p} + \sum_{k=0}^{K} \beta_k \cdot \varepsilon_{t-k} + u_t \tag{25}$$

where the $(\varepsilon_{t-k})_{k=0}^{K}$ are the identified shocks. The authors use information criteria to pick an optimal lag length combination (P, K). In the empirical application, for estimating the response of inflation, forecasts, and forecast errors to the technology shocks, they find that $K = 1$ and $P = 1$ uniformly fits the data the best

[34] There are two salient differences. The first is that Coibion and Gorodnichenko (2012) use forecast means rather than medians as a measure of the aggregate. The second is that Coibion and Gorodnichenko (2012) measure expected annual inflation with the forecast of the 4-quarter-ahead price level relative to the now-cast of the (unreleased) current-quarter price level; whereas our main analysis uses three-quarters-ahead relative to the previous quarter.

subject to their chosen penalty for extra parameters. Finally, given the empirical ARMA representation, they can directly compute impulse response coefficients.

Our approach. This paper's approach in Section 4 is similar but has two key differences. First, we fix a larger value of K (in our preferred specification, $K = 3$), in anticipation of the fact that the model may demand more complex dynamics than an AR(1). Second, we instrument for lagged values of x_t using past shocks. Intuitively, this is meant to isolate the possibility that dynamics may be "shock-specific" and not informed entirely by the unconditional auto-covariance patterns in x_t. More formally, this is to be expected if the data-generating process does in fact involve multiple shocks and/or variables, so thinking of the model as exactly a single-shock ARMA could be very inaccurate.

For comparability with (25), we will estimate the following system of equations with two-stage least squares. The reduced-form equation is exactly (25) with $K = 1$ and $P = 3$ (to capture higher-order dynamics):

$$x_t = \alpha + \sum_{p=1}^{3} \gamma_p x_{t-p} + \sum_{k=0}^{1} \beta_k \varepsilon_{t-k} + u_t \tag{26}$$

The "first-stage" relates the lags of x_t with shocks before $t-1$. In vector form, if $X_{t-1} := [x_{t-1}, x_{t-2}, x_{t-3}]$ and $\mathscr{E}_{t-2} = [\varepsilon_{t-2-j}]_{\ell=0}^{J}$(with $J = 8$, like in the main text), has the following form:

$$X_{t-1} = \eta + \mathscr{E}'_{t-2}\Theta + e_t \tag{27}$$

Armed with these IV estimates of the γ and β coefficients, we can calculate an alternative impulse response.

Local projections. Finally, we can also run the following local projection regression separately for each horizon h:

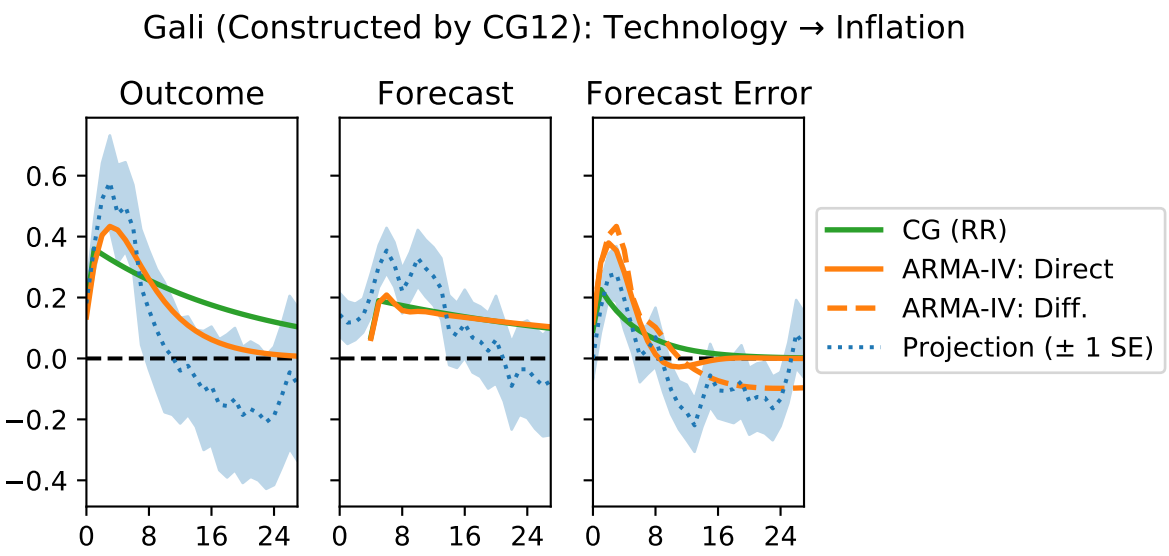
$$x_{t+h} = \alpha_h + \beta_{h,d} \cdot \varepsilon_t + \gamma' W_t + u_{t+h} \tag{28}$$

For controls W_t we will use the four lags each of labor productivity, the change in labor hours, and inflation that entered the original VAR. This is necessary, in the smaller sample, to make the estimated shock series truly orthogonal to lagged macro conditions.

B.2 Results

Figure A.6 compares the results, extended out to 28 quarters. Plotted in the blue dotted line, with a shaded 68% confidence interval, is the projection estimate of impulse responses for outcomes (left), forecasts (middle), and forecast errors. Plotted in green is the point estimate of the Coibion and Gorodnichenko (2012) method, or the estimate that comes from (25). Plotted in orange are the estimates from the IV method, or the combination of (26) and (27). And plotted in the orange dashed line is the difference between the orange lines for outcomes and forecasts, which is a different estimator for the response of forecast errors.

Figure A.6: Comparison of IRF Methods for Response to Technology Shock



The green lines in all cases are much more persistent than the projection responses. In the first and third case, in particular, the green lines smoothly and slowly converge back to zero. The unrestricted projection estimator, however, suggests that the response of inflation eventually turns negative (slightly, but not completely, offsetting the effects on the price level) and that the response of forecast errors also turns negative.

The ARMA-IV estimator, compared to the Coibion and Gorodnichenko (2012) method, gives a very similar response of forecast errors but a much less persistent response of the outcome. This estimation of the outcome IRF more closely matches the projection estimates. As such the "difference" estimator, or the dashed orange line in the third panel, shows evidence of over-extrapolation in the point estimate at moderate (>10 quarter) horizons. The ARMA-IV estimator directly applied to forecast errors, on the other hand, shows only modest evidence of over-shooting.

B.3 Suggestions for practice

The upshot of this might be summarized in the following points:

1. A method that imposes uniform dynamics as if the data-generating process involved only one shock, like that introduced by Romer and Romer (2004) and adopted by Coibion and Gorodnichenko (2012), may provide a particularly incomplete picture of dynamics of forecasts and forecast errors.

2. The possible solutions include the "shock-specific" IV approach introduced here, a flexible local-projection, or a more structured multi-variate model. The trade-offs between these models involve robustness and small-sample efficiency.

3. For the first and third methods discussed above, it is not unreasonable based on the theory to have an informed prior to favor dynamics that are more complicated than an AR(1) or ARMA(1,K) to char-

Table A.6: Predicting Individual Forecast Errors with Aggregate Revisions (All Horizons)

$$\text{Error}_{i,t,k} = -K_{noise} \cdot (\text{Revision}_{i,t,k} - \text{Revision}_{t,k}) + K_{agg} \cdot \text{Revision}_{t,k} + \alpha + u_{i,t,k}$$

	(1)	(2)	(3)	(4)	(5)	(6)
	Unemployment					
horizon k	1		2		3	
Period	1968-2017	1984-2017	1968-2017	1984-2017	1968-2017	1984-2017
$\text{Revision}_{i,t,k}$-$\text{Revision}_{t,k}$ ($-K_{noise}$)	-0.183	-0.217	-0.189	-0.264	-0.166	-0.162
	(0.035)	(0.039)	(0.043)	(0.043)	(0.043)	(0.053)
$\text{Revision}_{t,k}$ (K_{agg})	0.441	0.462	0.642	0.722	0.745	0.841
	(0.114)	(0.159)	(0.138)	(0.183)	(0.173)	(0.210)
R^2	0.120	0.136	0.147	0.195	0.103	0.152
Observations	5808	3986	5699	3918	5383	3769

	(7)	(8)	(9)	(10)	(11)	(12)
	Inflation					
horizon k	1		2		3	
Period	1968-2017	1984-2017	1968-2017	1984-2017	1968-2017	1984-2017
$\text{Revision}_{i,t,k}$-$\text{Revision}_{t,k}$ ($-K_{noise}$)	-0.422	-0.517	-0.427	-0.481	-0.346	-0.410
	(0.047)	(0.034)	(0.036)	(0.035)	(0.042)	(0.041)
$\text{Revision}_{t,k}$ (K_{agg})	0.675	-0.070	1.108	0.179	1.550	0.412
	(0.209)	(0.185)	(0.245)	(0.178)	(0.278)	(0.180)
R^2	0.168	0.106	0.194	0.085	0.211	0.072
Observations	5496	3779	5458	3745	5147	3643

Notes: The observation is a forecaster by quarter between Q4-1968 and Q4-2017. Standard errors are clustered two-way by forecaster ID and time period. Both errors and revisions are winsorized over the sample to restrict to 4 times the inter-quartile range away from the median. The data used for outcomes are first-release ("vintage").

acterize forecasts and forecast errors. These more complicated dynamics open up the possibility for "over-shooting" beyond a (usually small) period K.

C Noise and a Hybrid Regression

Corollary 2 underscored that, away from rational expectations, the CG regression coefficient is no more a measure of the informational friction alone: it is "contaminated" by the departure from the rationality. But the BGMS coefficient isolates the role of the latter. This suggests that the *gap* between the two coefficients ought to say something about the actual level of noise. Condition (16) makes clear that this intuition is correct up to a rescaling: the term $\frac{\kappa_1 \tau^{-1}}{\gamma_{agg}}$, which isolates the effect of the *actual* noise, can be obtained by subtracting $\frac{\gamma_{ind}}{\gamma_{agg}} K_{BGMS}$ from K_{CG}.

We next show how one can arrive at essentially the same answer with a "hybrid" of the CG and BGMS regressions:Define the new coefficient

$$K_{noise} \equiv \frac{\kappa_1 \tau^{-1}}{\gamma_{agg}} = K_{CG} - \frac{\gamma_{ind}}{\gamma_{agg}} K_{BGMS} \tag{29}$$

In the theory, the following regression holds:

$$\text{Error}_{i,t,k} = K_{\text{CG}} \cdot \text{Revision}_{t,k} - K_{\text{noise}} \cdot \frac{\mathcal{V}_{\text{agg}}}{\mathcal{V}_{\text{idio}}} \cdot \left(\text{Revision}_{i,t,k} - \text{Revision}_{t,k}\right) + u_{i,t,k} \tag{30}$$

Furthermore, K_{noise} is non-negative and strictly increasing in the level of noise, or decreasing in τ, with $K_{\text{noise}} \to 0$ as $\tau \to \infty$.

From the perspective of this regression, K_{CG} measures the predictability in *individual* forecast errors attributed to the *common* component of the lagged forecast revisions, and K_{noise} the one attributed to the purely *idiosyncratic* components of the lagged forecast revisions. As already explained, the former confounds the effects of misspecification and information. The latter, which is again the gap between K_{CG} and K_{BGMS} appropriately rescaled, isolates the effect of the idiosyncratic noise.[35]

Table A.6 shows results from estimating the hybrid regression over the full and restricted samples for all horizons of forecast. Across these margins, the estimated value of K_{noise} is positive (and statistically different from zero). This is lines up with the following observation: if we go back to the results presented in Section 3 and the Appendix regarding Facts 1 and 2, we can readily verify that K_{BGMS} was consistently lower that K_{CG}, even in specifications where both were positive.[36]

Of course, as evident from Proposition C, the hybrid regression does not provide *independent* information compared to Facts 1 and 2: the coefficients of the hybrid regression can be inferred from the original CG and BGMS regressions, and vice versa.[37] What this regression however accomplishes is to *combine* Facts 1 and 2 in way that more clearly illustrates how the gap between K_{CG} and K_{BGMS}, or more precisely the object K_{noise} described above, provides the needed "correction" of the original CG coefficient. With rational expectations, K_{noise} coincides with K_{CG}. Away from that benchmark, K_{noise} partials out from K_{CG} the component due to irrationality. In both cases, K_{noise} isolates the effect of idiosyncratic noise.[38]

Let us close this detour with the following remark. So far, we have have focused on how the three frictions jointly shape K_{CG} and K_{BGMS}, or equivalently K_{CG} and K_{noise}, and how one could get a measure of the true level of noise. But as anticipated in the Introduction, the true level of noise is ultimately irrelevant for the particular question of how all three frictions combined affect the aggregate dynamics. In the sequel, we will therefore show how one can bypass the individual-level evidence, or Fact 2 and the coefficients K_{BGMS} and K_{noise}, and identify the overall effect on the aggregate dynamics solely from Facts 1 and 3, or from K_{CG} and the reversal of sign in the IRF of aggregate forecast errors.

[35] To the best of our knowledge, the particular regression we propose here and the offered structural interpretation are novel. However, Fuhrer (2018) and Kohlhas and Broer (2019) contain a few empirical specifications that have a similar spirit, namely the separately test the extent to which aggregate-level and ind individual-level variables help predict forecast errors.

[36] The same seems to be true for almost all the specifications considered in Bordalo et al. (2018), including those regarding a variety of interest rates and spreads.

[37] To be precise, one also needs to compute $\mathcal{V}_{\text{idio}}$ and $\mathcal{V}_{\text{agg}}$, the variances of, respectively, the individual and aggregate forecast revisions. But these variances are already implicit in the calculation of K_{BGMS} and K_{CG}: they are the variances of the respective regressors.

[38] The following caveat applies to the adopted interpretation of K_{noise}. In the model we work with in this paper, idiosyncratic noise is the sole source of heterogeneity in beliefs: irrationality is a (possibly time-varying) fixed effect in the cross-section of the population. Without this restriction, K_{noise} may confound the effects of "rational" noise (due to idiosyncratic information) and "irrational" noise (due to idiosyncratic misspecification).

D Proofs

Proof of Lemma 1

We normalize the scaling coefficient $r = 1$. The perceived signal process can be represented as

$$s_{i,t} = \mathbf{M}(\mathbb{L}) \begin{bmatrix} \eta_t \\ u_{i,t} \end{bmatrix}, \qquad \text{with} \qquad M(\mathbb{L}) = \begin{bmatrix} \frac{1}{1-\hat{\rho}\mathbb{L}} & \hat{\tau}^{-\frac{1}{2}} \end{bmatrix}.$$

Let $B(\mathbb{L})$ denote the fundamental representation of the perceived signal process,[39] which is given by

$$B(\mathbb{L}) = \hat{\tau}^{-\frac{1}{2}} \sqrt{\frac{\hat{\rho}}{\hat{\lambda}}} \frac{1-\hat{\lambda}\mathbb{L}}{1-\hat{\rho}\mathbb{L}}, \text{ where } \hat{\lambda} = \frac{1}{2}\left(\hat{\rho} + \frac{1+\hat{\tau}}{\hat{\rho}} + \sqrt{\left(\hat{\rho} + \frac{1+\hat{\tau}}{\hat{\rho}}\right)^2 - 4}\right).$$

It is useful to note that $\hat{\lambda} < \hat{\rho}$. By the Wiener-Hopf prediction formula, the individual forecast about z_t is

$$\mathbb{E}_{i,t}[z_t] = \left[\frac{1}{1-\hat{\rho}\mathbb{L}} M'(\mathbb{L}^{-1}) B(\mathbb{L}^{-1})^{-1}\right]_+ B(\mathbb{L})^{-1} s_{i,t} = \left(1 - \frac{\hat{\lambda}}{\hat{\rho}}\right) \frac{1}{1-\hat{\lambda}\mathbb{L}} s_{i,t}.$$

Alternatively, this forecast rule can be written as

$$\mathbb{E}_{i,t}[z_t] = (1-\hat{g})\hat{\rho}\,\mathbb{E}_{i,t-1}[z_{t-1}] + \hat{g}\, s_{i,t},$$

which is a weighted average of the prior $\rho \mathbb{E}_{it-1}[z_{t-1}]$ and the new signal $s_{i,t}$, where the weight on the signal is the Kalman gain $\hat{g} = 1 - \frac{\hat{\lambda}}{\hat{\rho}}$. In the equations above, note that only perceived $\hat{\rho}$ and $\hat{\tau}$ matter for how agents use their signals. The actual ρ and τ matter for how the signal $s_{i,t}$ evolves overtime.

Accordingly, the one-period ahead forecast is

$$\mathbb{E}_{i,t}[z_{t+1}] = \hat{\rho}\mathbb{E}_{i,t}[z_t] = (\hat{\rho} - \hat{\lambda}) \frac{1}{1-\hat{\lambda}\mathbb{L}} s_{i,t} = (\hat{\rho} - \hat{\lambda}) \frac{1}{1-\hat{\lambda}\mathbb{L}} \left(\frac{1}{1-\rho\mathbb{L}} \varepsilon_t + \tau^{-\frac{1}{2}} u_{i,t}\right).$$

The individual forecast error and revision are then straightforward to obtain:

$$\text{Error}_{i,t} = z_{t+1} - \mathbb{E}_{i,t}[z_{t+1}] = \frac{1-\hat{\rho}L}{(1-\rho L)(1-\hat{\lambda}L)} \varepsilon_{t+1} - \frac{\hat{\rho}-\hat{\lambda}}{1-\hat{\lambda}\mathbb{L}} \tau^{-\frac{1}{2}} u_{i,t},$$

$$\text{Revision}_{i,t} = \mathbb{E}_{i,t}[z_{t+1}] - \mathbb{E}_{i,t-1}[z_{t+1}] = \frac{(\hat{\rho}-\hat{\lambda})(1-\hat{\rho}L)}{(1-\rho L)(1-\hat{\lambda}L)} \varepsilon_t + \frac{(\hat{\rho}-\hat{\lambda})(1-\hat{\rho}L)}{1-\hat{\lambda}\mathbb{L}} \tau^{-\frac{1}{2}} u_{i,t}.$$

Proof of Proposition 1

First consider the calculation of K_{BGMS}.

$$\begin{aligned}
&\text{Cov}(\text{Error}_{i,t}, \text{Revision}_{i,t}) \\
=&\text{Cov}\left(\frac{1-\hat{\rho}\mathbb{L}}{(1-\rho\mathbb{L})(1-\hat{\lambda}\mathbb{L})} \varepsilon_{t+1}, \frac{(\hat{\rho}-\hat{\lambda})(1-\hat{\rho}\mathbb{L})}{(1-\rho\mathbb{L})(1-\hat{\lambda}\mathbb{L})} \varepsilon_t\right) + \text{Cov}\left(-\frac{\hat{\rho}-\hat{\lambda}}{1-\hat{\lambda}\mathbb{L}} \tau^{-\frac{1}{2}} u_{i,t}, \frac{(\hat{\rho}-\hat{\lambda})(1-\hat{\rho}\mathbb{L})}{1-\hat{\lambda}\mathbb{L}} \tau^{-\frac{1}{2}} u_{i,t}\right) \\
=&-(\hat{\rho}-\hat{\lambda}) \frac{\hat{\lambda}}{(1-\hat{\lambda}^2)} \hat{\tau}\left(\tau^{-1} - \hat{\tau}^{-1}\right) + (\rho-\hat{\rho})(\hat{\rho}-\hat{\lambda}) \frac{(1+\hat{\lambda}^2)(1-\rho^2) + (\hat{\lambda}+\rho)(\rho-\hat{\rho})}{(1-\hat{\lambda}^2)(1-\rho^2)(1-\hat{\lambda}\rho)}
\end{aligned}$$

[39] $B(\mathbb{L})$ satisfies the requirement $B(\mathbb{L})B'(\mathbb{L}^{-1}) = M(\mathbb{L})M'(\mathbb{L}^{-1})$ and $B(\mathbb{L})$ is invertible.

Denote κ_1 and κ_2 as

$$\kappa_1 = (\hat{\rho} - \hat{\lambda}) \frac{\hat{\lambda}}{(1-\hat{\lambda}^2)} \hat{\tau}, \quad \kappa_2 = (\hat{\rho} - \hat{\lambda}) \frac{(1+\hat{\lambda}^2)(1-\rho^2) + (\hat{\lambda}+\rho)(\rho-\hat{\rho})}{(1-\hat{\lambda}^2)(1-\rho^2)(1-\hat{\lambda}\rho)}.$$

It follows that

$$K_{BGMS} = \frac{-\kappa_1(\tau^{-1} - \hat{\tau}^{-1}) + \kappa_2(\rho - \hat{\rho})}{\mathcal{V}_{\text{ind}}}$$

As $1 > \hat{\rho} > \hat{\lambda} > 0$, $\kappa_1 > 0$. To show that $\kappa_2 > 0$, it is equivalent to show that $(1+\hat{\lambda}^2)(1-\rho^2) + (\hat{\lambda}+\rho)(\rho-\hat{\rho}) > 0$. Given that $\hat{\rho} < 1$, it follows that

$$(1+\hat{\lambda}^2)(1-\rho^2) + (\hat{\lambda}+\rho)(\rho-\hat{\rho}) > (1+\hat{\lambda}^2)(1-\rho^2) + (\hat{\lambda}+\rho)(\rho-1) = (1-\rho)(1-\hat{\lambda}+\hat{\lambda}^2(1+\rho)) > 0.$$

Now turn to the calculation of K_{CG}.

$$\begin{aligned}
\text{Cov}(\text{Error}_t, \text{Revision}_t) =& \text{Cov}\left(\frac{1-\hat{\rho}L}{(1-\rho L)(1-\hat{\lambda}L)} \varepsilon_{t+1}, \frac{(\hat{\rho}-\hat{\lambda})(1-\hat{\rho}L)}{(1-\rho L)(1-\hat{\lambda}L)} \varepsilon_t \right) \\
=& (\hat{\rho}-\hat{\lambda}) \left(\frac{\hat{\lambda}}{1-\hat{\lambda}^2} + (\rho-\hat{\rho}) \frac{(1+\hat{\lambda}^2)(1-\rho^2) + (\hat{\lambda}+\rho)(\rho-\hat{\rho})}{(1-\hat{\lambda}^2)(1-\rho^2)(1-\hat{\lambda}\rho)} \right) \\
=& \text{Cov}(\text{Error}_{i,t}, \text{Revision}_{i,t}) + \kappa_1 \tau^{-1},
\end{aligned}$$

which leads to

$$K_{CG} = \frac{\kappa_1 \tau^{-1}}{\mathcal{V}_{\text{agg}}} + \frac{\mathcal{V}_{\text{ind}}}{\mathcal{V}_{\text{agg}}} K_{\text{BGMS}}$$

Proof of Proposition C

We consider the case with $k = 1$. Note that average revision, Revision_t, and the idiosyncratic component of individual revision, $(\text{Revision}_{i,t} - \text{Revision}_t)$, are independent of each other. Therefore, the regression coefficient on the average forecast revision remains to be K_{CG}.

The covariance between individual forecast error and idiosyncratic revision component is

$$\begin{aligned}
\text{Cov}(\text{Error}_{i,t}, \text{Revision}_{i,t} - \text{Revision}_{i,t}) &= \text{Cov}\left(-\frac{\hat{\rho}-\hat{\lambda}}{1-\hat{\lambda}L} u_{it}, \frac{\hat{\rho}-\hat{\lambda}}{1-\hat{\lambda}L} (\tau^{-\frac{1}{2}} u_{i,t} - \hat{\rho}\tau^{-\frac{1}{2}} u_{i,t-1}) \right) \\
&= -(\hat{\rho}-\hat{\lambda}) \frac{\hat{\tau}}{\tau} \frac{\hat{\lambda}}{1-\hat{\lambda}^2} \\
&= -\kappa_1 \tau^{-1}.
\end{aligned}$$

Denote the regression coefficient on $(\text{Revision}_{i,t} - \text{Revision}_t)$ as β. It follows that

$$\begin{aligned}
\beta &= \frac{\text{Cov}(\text{Error}_{i,t}, \text{Revision}_{i,t} - \text{Revision}_t)}{\mathcal{V}_{\text{idio}}} = \frac{\text{Cov}(\text{Error}_{i,t}, \text{Revision}_{i,t}) - \text{Cov}(\text{Error}_t, \text{Revision}_t)}{\mathcal{V}_{\text{idio}}} \\
&= \frac{\mathcal{V}_{\text{ind}}}{\mathcal{V}_{\text{idio}}} K_{\text{BGMS}} - \frac{\mathcal{V}_{\text{agg}}}{\mathcal{V}_{\text{idio}}} K_{\text{CG}},
\end{aligned}$$

and

$$K_{\text{noise}} \equiv \frac{\kappa_1 \tau^{-1}}{\mathcal{V}_{\text{agg}}} = -\beta \frac{\mathcal{V}_{\text{idio}}}{\mathcal{V}_{\text{agg}}} = K_{\text{CG}} - \frac{\mathcal{V}_{\text{ind}}}{\mathcal{V}_{\text{agg}}} K_{\text{BGMS}}$$

Because κ_1 and $\mathcal{V}_{\text{agg}}$ are independent of τ, K_{noise} is decreasing in τ, and vanishes when $\tau \to \infty$.

Proof of Proposition 2

The law of motion of the average forecast error is given by

$$\text{Error}_t = \frac{1-\hat{\rho}\mathbb{L}}{(1-\rho\mathbb{L})(1-\hat{\lambda}\mathbb{L})}\varepsilon_{t+1} = \left(\frac{\rho-\hat{\rho}}{\rho-\hat{\lambda}}\frac{1}{1-\rho L} + \frac{\hat{\rho}-\hat{\lambda}}{\rho-\hat{\lambda}}\frac{1}{1-\hat{\lambda}L}\right)\varepsilon_{t+1}.$$

Suppose $\rho > \hat{\rho}$, then $\rho > \hat{\lambda}$. The coefficients of the two AR(1) terms are both positive, and the responses are therefore all positive.

Suppose $\rho < \hat{\rho}$. Consider the following continuous time version of the response

$$g(t) = \frac{\rho-\hat{\rho}}{\rho-\hat{\lambda}}\rho^t + \frac{\hat{\rho}-\hat{\lambda}}{\rho-\hat{\lambda}}\hat{\lambda}^t,$$

and $g(t) = \zeta_k$ when $t = k \in \{0,1,\ldots\}$. Note that: (1) $g(t)$ is negative when t is large enough (no matter $\rho > \hat{\lambda}$ or $\rho < \hat{\lambda}$); (2) when $t = 0$, $g(0) = 1 > 0$; (3) there is at most one root of $g(t)$. As a result, $\{\zeta_k\}_{k=1}^{\infty}$ eventually stay negative, but they might be positive or negative for k small enough.

The root of $g(t)$ is

$$K_{\text{IRF}} = \frac{\log(\hat{\rho}-\rho) - \log(\hat{\rho}-\hat{\lambda})}{\log\hat{\lambda} - \log\rho}.$$

To have $\{\zeta_k\}_{k=1}^{\infty}$ switch signs, it is necessary that $g(1) > 0$ and $\hat{\rho} > \rho$, which correspond to

$$g(1) = \rho + \hat{\lambda} - \hat{\rho} > 0, \quad \text{and} \quad \hat{\rho} > \rho,$$

or

$$\hat{\lambda} > \hat{\rho} - \rho, \quad \text{and} \quad \hat{\rho} > \rho.$$

When $\hat{\rho} > \rho$ but $\hat{\lambda} > \hat{\rho} - \rho$, the sequences $\{\zeta_k\}_{k=1}^{\infty}$ stay negative all the time.

Proof of Proposition 3

The aggregate consumption satisfy the fixed point restriction

$$c_t = \sum_{k=0}^{\infty}\beta^k\bar{\mathbb{E}}_t[\xi_{t+k}] + (1-\beta)\sum_{k=0}^{\infty}\beta^k\bar{\mathbb{E}}_t[c_{t+k+1}],$$

where we have used the market clearing condition $y_t = c_t$, and the assumption that agents observe y_t but not extract information from it. This aggregate outcome is the outcome of the following beauty-contest game

$$c_{i,t} = \mathbb{E}_{i,t}[\xi_t] + \beta\mathbb{E}_{i,t}[c_{i,t+1}] + (1-\beta)\mathbb{E}_{i,t}[c_{t+1}].$$

Denote the agent's equilibrium policy function as

$$c_{i,t} = h(L)s_{i,t}$$

for some lag polynomial $h(\mathbb{L})$. The actual law of motion of aggregate outcome can then be expressed as follows

$$c_t = h(\mathbb{L})\xi_t = \frac{h(\mathbb{L})}{1-\rho\mathbb{L}}\varepsilon_t.$$

However, the perceived law of motion by consumers is

$$c_t = \frac{h(\mathbb{L})}{1-\hat{\rho}\mathbb{L}}\varepsilon_t.$$

Similar to the case where the outcome is given by the exogenous AR(1) process, the forecast about the fundamental is

$$\mathbb{E}_{i,t}[\xi_t] = \left(1 - \frac{\hat{\lambda}}{\hat{\rho}}\right) \frac{1}{1-\hat{\lambda}\mathbb{L}} s_{i,t} \equiv G_1(\mathbb{L}) s_{i,t}.$$

Consider the forecast of the future own and average actions. The perceived law of motion of $c_{i,t+1}$ and c_{t+1} are

$$c_{t+1} = \begin{bmatrix} \frac{h(\mathbb{L})}{L(1-\hat{\rho}\mathbb{L})} & 0 \end{bmatrix} \begin{bmatrix} \varepsilon_t \\ u_{i,t} \end{bmatrix}, \quad c_{i,t+1} - c_{t+1} = \begin{bmatrix} 0 & \hat{\tau}^{-\frac{1}{2}} \frac{h(\mathbb{L})}{\mathbb{L}} \end{bmatrix} \begin{bmatrix} \varepsilon_t \\ u_{i,t} \end{bmatrix},$$

and the forecasts are

$$\mathbb{E}_{i,t}[c_{t+1}] = G_2(\mathbb{L}) s_{i,t}, \qquad G_2(\mathbb{L}) \equiv \frac{\hat{\lambda}}{\hat{\rho}} \hat{\tau} \left(\frac{h(\mathbb{L})}{(1-\hat{\lambda}\mathbb{L})(\mathbb{L}-\hat{\lambda})} - \frac{h(\hat{\lambda})(1-\hat{\rho}\mathbb{L})}{(1-\hat{\rho}\hat{\lambda})(\mathbb{L}-\hat{\lambda})(1-\hat{\lambda}\mathbb{L})} \right),$$

$$\mathbb{E}_{i,t}\left[c_{i,t+1} - c_{t+1}\right] = G_3(\mathbb{L}) s_{i,t}, \qquad G_3(\mathbb{L}) \equiv \frac{\hat{\lambda}}{\hat{\rho}} \left(\frac{h(\mathbb{L})(\mathbb{L}-\hat{\rho})}{\mathbb{L}(\mathbb{L}-\hat{\lambda})} - \frac{h(\hat{\lambda})(\hat{\lambda}-\hat{\rho})}{\hat{\lambda}(\mathbb{L}-\hat{\lambda})} - \frac{\hat{\rho}}{\hat{\lambda}} \frac{h(0)}{\mathbb{L}} \right) \frac{1-\hat{\rho}\mathbb{L}}{1-\hat{\lambda}\mathbb{L}}$$

Recall that fixed point problem that characterizes the equilibrium is

$$c_{i,t} = \mathbb{E}_{i,t}[\xi_t] + \beta \mathbb{E}_{i,t}[c_{i,t+1}] + (1-\beta)\mathbb{E}_{i,t}[c_{t+1}].$$

We can replace the left-hand side with $h(L)s_{i,t}$. Using the results derived above, on the other hand, we can replace the right-hand side with $\left[G_1(L) + G_2(L) + \beta G_3(L)\right] s_{i,t}$. It follows that in equilibrium

$$h(L) = G_1(L) + G_2(L) + \beta G_3(L).$$

Equivalently, we need to find an analytic function $h(z)$ that solves

$$\begin{aligned} h(z) = {} & \frac{\hat{\lambda}}{\hat{\rho}} \hat{\tau} \frac{1}{1-\hat{\rho}\hat{\lambda}} \frac{1}{1-\hat{\lambda}z} + \frac{\hat{\lambda}}{\hat{\rho}} \hat{\tau} \left(\frac{h(z)}{(1-\hat{\lambda}z)(z-\hat{\lambda})} - \frac{h(\hat{\lambda})(1-\hat{\rho}z)}{(1-\hat{\rho}\hat{\lambda})(z-\hat{\lambda})(1-\hat{\lambda}z)} \right) \\ & + \beta \frac{\hat{\lambda}}{\hat{\rho}} \left(\frac{h(z)(z-\hat{\rho})}{z(z-\hat{\lambda})} - \frac{h(\hat{\lambda})(\hat{\lambda}-\hat{\rho})}{\hat{\lambda}(z-\hat{\lambda})} - \frac{\hat{\rho}}{\hat{\lambda}} \frac{h(0)}{z} \right) \frac{1-\hat{\rho}z}{1-\hat{\lambda}z}, \end{aligned}$$

which can be transformed as

$$\tilde{C}(z)h(z) = d(z; h(\hat{\lambda}), h(0))$$

where

$$\begin{aligned} \tilde{C}(z) &\equiv z(1-\hat{\lambda}z)(z-\hat{\lambda}) - \frac{\hat{\lambda}}{\hat{\rho}} \left\{ \beta(z-\hat{\rho})(1-\hat{\rho}z) + \hat{\tau}z \right\} \\ d(z; h(\hat{\lambda}), h(0)) &\equiv \frac{\hat{\lambda}}{\hat{\rho}} \hat{\tau} \frac{1}{1-\hat{\rho}\hat{\lambda}} z(z-\hat{\lambda}) - \frac{1}{\hat{\rho}} \left(\hat{\tau} \frac{\hat{\lambda}}{1-\hat{\rho}\hat{\lambda}} + \beta(\hat{\lambda}-\hat{\rho}) \right) z(1-\hat{\rho}z)h(\hat{\lambda}) - \beta(z-\hat{\lambda})(1-\hat{\rho}z)h(0) \end{aligned}$$

Note that $\tilde{C}(z)$ is a cubic equation and therefore contains with three roots. We will verify later that there are two inside roots and one outside root. To make sure that $h(z)$ is an analytic function, we choose $h(0)$ and $h(\hat{\lambda})$ so that the two roots of $d(z; h(\hat{\lambda}), h(0))$ are the same as the two inside roots of $\tilde{C}(z)$. This pins down the constants $\{h(0), h(\hat{\lambda})\}$, and therefore the policy function $h(L)$ is

$$h(\mathbb{L}) = \left(1 - \frac{\vartheta}{\hat{\rho}}\right) \frac{1}{1-\hat{\rho}} \frac{1}{1-\vartheta\mathbb{L}},$$

where ϑ^{-1} is the root of $\tilde{C}(z)$ outside the unit circle.

Now we verify that $\tilde{C}(z)$ has two inside roots and one outside root. $\tilde{C}(z)$ can be rewritten as $\hat{\lambda}C(z)$ where

$$\begin{aligned} C(z) &= -z^3 + \left(\hat{\rho} + \frac{1}{\hat{\rho}} + \frac{1}{\hat{\rho}}\hat{\tau} + \beta\right) z^2 - \left(1 + \beta\left(\hat{\rho} + \frac{1}{\hat{\rho}}\right) + \frac{1}{\hat{\rho}}\hat{\tau}\right) z + \beta, \\ &= -z^3 + \left(\hat{\rho} + \frac{1}{\hat{\rho}} + \frac{1}{\hat{\rho}}\hat{\tau} + 1 - \text{mpc}\right) z^2 - \left(1 + (1-\text{mpc})\left(\hat{\rho} + \frac{1}{\hat{\rho}}\right) + \frac{1}{\hat{\rho}}\hat{\tau}\right) z + 1 - \text{mpc}. \end{aligned}$$

With the assumption that $1 > \text{mpc} > 0$, it is straightforward to verify that the following properties hold:

$$C(0) = 1 - \text{mpc} > 0, \quad C(\hat{\lambda}) = -\text{mpc}\frac{\hat{\tau}}{\hat{\rho}} < 0, \quad C(1) = \text{mpc}\left(\frac{1}{\hat{\rho}} + \hat{\rho} - 2\right) > 0.$$

Therefore, the three roots are all real, two of them are between 0 and 1, and the third one ϑ^{-1} is larger than 1.

To show that ϑ is less than $\hat{\rho}$, it is sufficient to show that

$$C\left(\frac{1}{\hat{\rho}}\right) = \frac{\hat{\tau}(1-\hat{\rho})}{\hat{\rho}^3} > 0.$$

Since $C(\vartheta^{-1}) = 0$, it has to be that ϑ^{-1} is larger than $\hat{\rho}^{-1}$, or $\vartheta < \hat{\rho}$.

It also implies that $C(z)$ is decreasing in z in the neighborhood of $z = \vartheta^{-1}$, a property that we use to characterize comparative statics of ϑ. Taking derivative of $C(z)$ with respect to mpc, and evaluating that derivative at $z = \vartheta^{-1}$, we get

$$\frac{\partial C(\vartheta^{-1})}{\partial \text{mpc}} = -(\vartheta^{-1} - \hat{\rho})(\vartheta^{-1} - \hat{\rho}^{-1}) < 0$$

Combining this with the earlier observation that $\frac{\partial C(\vartheta^{-1})}{\partial z} < 0$, and using the Implicit Function Theorem, we infer that ϑ is an increasing function of mpc.

When mpc=0, we have

$$c_{i,t} = \mathbb{E}_{i,t}[\xi_t] + \mathbb{E}_{i,t}[c_{i,t+1}] = \sum_{k=0}^{\infty} \mathbb{E}_{i,t}[\xi_{t+k}] = \frac{1}{1-\rho}\mathbb{E}_{i,t}[\xi_t],$$

and

$$y_t = c_t = \frac{1}{1-\rho}\bar{\mathbb{E}}_t[\xi_t] = \frac{1}{1-\rho}\left(1 - \frac{\hat{\lambda}}{\hat{\rho}}\right)\frac{1}{1-\hat{\lambda}\mathbb{L}}\frac{1}{1-\rho\mathbb{L}}\varepsilon_t.$$

That is, $\vartheta = \hat{\lambda}$ in this case.

In Angeletos and Huo (2019), the equilibrium policy rule is derived under the assumption that $\rho = \hat{\rho}$ and $\tau = \hat{\tau}$. In the derivation above, note that $h(\mathbb{L})$ does not depend on ρ nor τ. The actual law of motion of $y_t = c_t$ will depend on ρ

$$y_t = \frac{1}{1-\hat{\rho}}\left(1 - \frac{\vartheta}{\hat{\rho}}\right)\frac{1}{1-\vartheta\mathbb{L}}\frac{1}{1-\rho\mathbb{L}}\varepsilon_t.$$

On the other hand, the frictionless case is given by

$$y_t^* = \frac{1}{1-\rho}\frac{1}{1-\rho\mathbb{L}}\varepsilon_t.$$

Combining these two leads to

$$y_t = \left(1 - \frac{\vartheta}{\hat{\rho}}\right)\left(1 + \frac{\hat{\rho}-\rho}{1-\hat{\rho}}\right)\left(\frac{1}{1-\vartheta\mathbb{L}}\right)y_t^*.$$

Turn to the forecast of the future outcome. By the Wiener-Hopf prediction formula, the individual forecast about y_{t+1} is

$$\begin{aligned} \mathbb{E}_{i,t}[y_{t+1}] &= \left[\frac{1}{1-\hat{\rho}}\left(1 - \frac{\vartheta}{\hat{\rho}}\right)\frac{1}{1-\vartheta\mathbb{L}}\frac{1}{1-\rho\mathbb{L}}M'(\mathbb{L}^{-1})B(\mathbb{L}^{-1})^{-1}\right]_+ B(\mathbb{L})^{-1}s_{i,t}, \\ &= \frac{1}{1-\hat{\rho}}\left(1 - \frac{\vartheta}{\hat{\rho}}\right)\left(1 - \frac{\hat{\lambda}}{\hat{\rho}}\right)\frac{1}{1-\vartheta\hat{\lambda}}\frac{\hat{\rho} + \vartheta - \hat{\rho}\vartheta(\mathbb{L} + \hat{\lambda})}{(1-\vartheta\mathbb{L})(1-\hat{\lambda}\mathbb{L})}s_{i,t}, \end{aligned}$$

and the average forecast is

$$\bar{\mathbb{E}}_t[y_{t+1}] = \left(1-\frac{\hat{\lambda}}{\hat{\rho}}\right)\frac{1}{1-\vartheta\hat{\lambda}}\frac{\hat{\rho}+\vartheta-\hat{\rho}\vartheta(\mathbb{L}+\hat{\lambda})}{(1-\vartheta\mathbb{L})(1-\hat{\lambda}\mathbb{L})}\left(1-\frac{\vartheta}{\hat{\rho}}\right)\left(1+\frac{\hat{\rho}-\rho}{1-\hat{\rho}}\right)y_t^*$$

Proof of Proposition 4

Denote $\kappa \equiv \left(1-\frac{\vartheta}{\hat{\rho}}\right)\left(1+\frac{\hat{\rho}-\rho}{1-\hat{\rho}}\right)\frac{1}{1-\rho}$. If $c_t = \kappa\frac{1}{(1-\vartheta L)(1-\rho\mathbb{L})}$ is the outcome of the perfect information outcome, it has to be that

$$\begin{aligned} c_t =& \xi_t + \omega_f\mathbb{E}_t^*[c_{t+1}] + \omega_b c_{t-1} \\ =& \frac{1}{1-\rho\mathbb{L}} + \omega_f\kappa\frac{\vartheta+\rho-\vartheta\rho\mathbb{L}}{(1-\vartheta\mathbb{L})(1-\rho\mathbb{L})} + \omega_b\kappa\frac{\mathbb{L}}{(1-\vartheta\mathbb{L})(1-\rho\mathbb{L})} \end{aligned}$$

where the right-hand side is simply the perfect information expectation of the behavioral equilibrium. This leads to

$$\begin{aligned} \omega_f &= \frac{\hat{\rho}^2-\vartheta}{(\vartheta+\rho)(\hat{\rho}-\vartheta)}, \\ \omega_b &= \frac{\vartheta(\rho(\hat{\rho}-\vartheta)+\vartheta\hat{\rho}(1-\hat{\rho}))}{(\vartheta+\rho)(\hat{\rho}-\vartheta)}. \end{aligned}$$

Proof of Proposition 6

With $\rho = \hat{\rho}$, it follows that

$$\begin{aligned} \omega_f &= \frac{\rho^2-\vartheta}{(\vartheta+\rho)(\rho-\vartheta)}, \\ \omega_b &= \frac{\vartheta\rho^2(1-\vartheta)}{(\vartheta+\rho)(\rho-\vartheta)}. \end{aligned}$$

Note that

$$\omega_f < \frac{\rho^2-\vartheta^2}{(\vartheta+\rho)(\rho-\vartheta)} = 1$$

and $\omega_b > 0$ as $\vartheta < \rho < 1$. For the comparative statics, we have

$$\begin{aligned} \frac{\partial\omega_f}{\partial\vartheta} &= \frac{-(\rho^2+\vartheta^2)+2\rho^2\vartheta}{(\rho^2-\vartheta^2)^2} < \frac{-(\rho^2+\vartheta)+2\rho\vartheta}{(\rho^2-\vartheta^2)^2} = \frac{-(\rho-\vartheta)^2}{(\rho^2-\vartheta^2)^2} < 0 \\ \frac{\partial\omega_b}{\partial\vartheta} &= \frac{\rho^2(\rho^2+\vartheta^2-2\vartheta\rho^2)}{(\rho^2-\vartheta^2)^2} > \frac{\rho^2(\rho^2+\vartheta^2-2\vartheta\rho)}{(\rho^2-\vartheta^2)^2} = \left(\frac{\rho}{\rho+\vartheta}\right)^2 > 0 \end{aligned}$$

Since ϑ is increasing in mpc, ω_f is decreasing in mpc and ω_b is increasing in mpc.

Now consider the regression coefficient K_{CG} and K_{BGMS}. The individual forecast error and forecast revision are given by

$$\begin{aligned} y_{t+1} - \mathbb{E}_{i,t}[y_{t+1}] &= \frac{1}{1-\rho}\left(1-\frac{\vartheta}{\rho}\right)\left(g_1^{\varepsilon}(\mathbb{L})\varepsilon_{t+1} + g_1^{u}(\mathbb{L})u_{i,t}\right), \\ \mathbb{E}_{i,t}[y_{t+1}] - \mathbb{E}_{i,t-1}[y_{t+1}] &= \frac{1}{1-\rho}\left(1-\frac{\vartheta}{\rho}\right)\left(g_2^{\varepsilon}(\mathbb{L})\varepsilon_{t} + g_2^{u}(\mathbb{L})u_{i,t}\right), \end{aligned}$$

where

$$g_1^{\varepsilon}(\mathbb{L}) = \frac{1}{(1-\vartheta\mathbb{L})(1-\rho\mathbb{L})} - \left(1-\frac{\hat{\lambda}}{\rho}\right)\frac{1}{1-\vartheta\hat{\lambda}}\frac{\rho+\vartheta-\rho\vartheta(\mathbb{L}+\hat{\lambda})}{(1-\vartheta\mathbb{L})(1-\hat{\lambda}\mathbb{L})(1-\rho\mathbb{L})}\mathbb{L},$$

$$g_2^{\varepsilon}(\mathbb{L}) = \left(1-\frac{\hat{\lambda}}{\rho}\right)\frac{1}{1-\vartheta\hat{\lambda}}\left(\frac{\rho+\vartheta-\rho\vartheta(\mathbb{L}+\hat{\lambda})}{(1-\vartheta\mathbb{L})(1-\hat{\lambda}\mathbb{L})(1-\rho\mathbb{L})}(1-(\vartheta+\rho)\mathbb{L}) + \frac{\rho\vartheta(1-\rho\vartheta\hat{\lambda}\mathbb{L})}{(1-\vartheta\mathbb{L})(1-\hat{\lambda}\mathbb{L})(1-\rho\mathbb{L})}\mathbb{L}\right),$$

$$g_1^{u}(\mathbb{L}) = -\left(1-\frac{\hat{\lambda}}{\rho}\right)\frac{1}{1-\vartheta\hat{\lambda}}\frac{\rho+\vartheta-\rho\vartheta(\mathbb{L}+\hat{\lambda})}{(1-\vartheta\mathbb{L})(1-\hat{\lambda}\mathbb{L})}\tau^{-1},$$

$$g_2^{u}(\mathbb{L}) = \left(1-\frac{\hat{\lambda}}{\rho}\right)\frac{1}{1-\vartheta\hat{\lambda}}\left(\frac{\rho+\vartheta-\rho\vartheta(\mathbb{L}+\hat{\lambda})}{(1-\vartheta\mathbb{L})(1-\hat{\lambda}\mathbb{L})}(1-(\vartheta+\rho)\mathbb{L}) + \frac{\rho\vartheta(1-\rho\vartheta\hat{\lambda}\mathbb{L})}{(1-\vartheta\mathbb{L})(1-\hat{\lambda}\mathbb{L})}\mathbb{L}\right)\tau^{-1}.$$

The covariance between individual forecast error and individual forecast revision is

$$\mathrm{Cov}\left(\mathrm{Error}_{i,t},\, \mathrm{Revision}_{i,t}\right) = \left(\frac{1}{1-\rho}\left(1-\frac{\vartheta}{\rho}\right)\right)^2 \left(\mathrm{Cov}\left(g_1^{\varepsilon}(\mathbb{L})\varepsilon_{t+1}, g_2^{\varepsilon}(\mathbb{L})\varepsilon_t\right) \;+\; \mathrm{Cov}\left(g_1^{u}(\mathbb{L})u_{i,t}, g_2^{u}(\mathbb{L})u_{i,t}\right)\right),$$

and a long but straightforward calculation yields the following expression:

$$\mathrm{Cov}(g_1^{u}(\mathbb{L})u_{i,t}, g_2^{u}(\mathbb{L})u_{i,t}) = -\tau^{-1}\left(\left(1-\frac{\hat{\lambda}}{\rho}\right)\frac{1}{1-\vartheta\hat{\lambda}}\right)^2 \frac{1-\hat{\lambda}\rho}{(1-\hat{\lambda}\vartheta)(1-\hat{\lambda}^2)}\Delta,$$

where

$$\Delta \equiv \left(\vartheta^3\hat{\lambda}\left(1-\hat{\lambda}^2\right) - 3\vartheta\hat{\lambda}\left(1-\vartheta\hat{\lambda}\right) + \left(1-\vartheta^2\right)\right)\rho^2 - \left(\vartheta^3\left(1-\hat{\lambda}^2\right) + \vartheta\left(3\vartheta\hat{\lambda}-2\right)\right)\rho + \vartheta^2.$$

We will verify that $\Delta > 0$ at the end of this proof.

If $\tau = \hat{\tau}$, then agents are rational and $K_{\mathrm{BGMS}} = 0$, that is, $\mathrm{Cov}\left(\mathrm{Error}_{i,t},\, \mathrm{Revision}_{i,t}\right) = 0$. It follows that

$$\begin{aligned}\mathrm{Cov}\left(\mathrm{Error}_t,\, \mathrm{Revision}_t\right) &= \left(\frac{1}{1-\rho}\left(1-\frac{\vartheta}{\rho}\right)\right)^2 \mathrm{Cov}\left(g_1^{\varepsilon}(\mathbb{L})\varepsilon_{t+1}, g_2^{\varepsilon}(\mathbb{L})\varepsilon_t\right)\\ &= -\left(\frac{1}{1-\rho}\left(1-\frac{\vartheta}{\rho}\right)\right)^2 \mathrm{Cov}\left(g_1^{u}(\mathbb{L})u_{i,t}, g_2^{u}(\mathbb{L})u_{i,t}\right) > 0,\end{aligned}$$

which implies that $K_{\mathrm{CG}} > 0$. If we fix the perceived $\hat{\tau}$ and vary the actual τ, it will not affect the average forecast error and forecast revision. Therefore, $K_{\mathrm{CG}} > 0$ even with $\tau \neq \hat{\tau}$.

However, when we fix the perceived $\hat{\tau}$ and vary the actual τ, this will change $\mathrm{Cov}(g_1^{u}(\mathbb{L})u_{i,t}, g_2^{u}(\mathbb{L})u_{i,t})$. To sign K_{BGMS} away from the $\tau = \hat{\tau}$ benchmark, it is sufficient to check whether $\mathrm{Cov}(g_1^{u}(\mathbb{L})u_{i,t}, g_2^{u}(\mathbb{L})u_{i,t})$ is increasing or decreasing in τ. It follows that $\mathrm{Cov}(g_1^{u}(\mathbb{L})u_{i,t}, g_2^{u}(\mathbb{L})u_{i,t})$ is strictly increasing in τ and, therefore, K_{BGMS} switches sign from negative to positive as τ crosses $\hat{\tau}$ from below. That is,

$$\mathrm{sign}\left(K_{BGMS}\right) = \mathrm{sign}\left(\tau - \hat{\tau}\right).$$

The argument is completed by the lemma below, which helps verify that $\Delta > 0$ by mapping ρ to x, ϑ to y, and λ to z.

Lemma. When $x, y, z \in (0,1)$, the following inequality holds

$$\left(y^3 z\left(1-z^2\right) - 3yz\left(1-yz\right) + \left(1-y^2\right)\right)x^2 - \left(y^3\left(1-z^2\right) + y\left(3yz-2\right)\right)x + y^2 > 0$$

Proof. Recast the left hand side of the above inequality as a quadratic in x:

$$C(x) \equiv \left(y^3 z\left(1-z^2\right) - 3yz\left(1-yz\right) + \left(1-y^2\right)\right)x^2 - \left(y^3\left(1-z^2\right) + y\left(3yz-2\right)\right)x + y^2.$$

This has two real roots, $x = x_1$ and $x = x_2$, given by

$$x_1 = -\frac{y}{1-yz} \qquad \text{and} \qquad x_2 = -\frac{y}{y^2z^2 - 2yz - y^2 + 1}.$$

Clearly, given the assumption that $y, z \in (0,1)$, x_1 is negative and $C(0) = y^2 > 0$. If x_2 is negative, then $C(x)$ is positive when $x \in (0,1)$. If x_2 is positive, to guarantee that $C(x)$ is positive when $x \in (0,1)$, we need to show that $x_2 > 1$, which is equivalent to show that

$$y^2z^2 - 2yz + \left(y - y^2 + 1\right) > 0$$

Define the following quadratic equation in z:

$$D(z) = y^2z^2 - 2yz + \left(y - y^2 + 1\right).$$

Its discriminant is $-4y^3\left(1-y\right)$, which is negative given that $y \in (0,1)$. Therefore, $D(z)$ is always positive, which in turn verifies $x_2 > 1$.

□

Lastly, the IRF of the average forecast error is

$$\begin{aligned}
y_{t+1} - \bar{\mathbb{E}}_t[y_{t+1}] &= \frac{1}{1-\rho}\left(1-\frac{\vartheta}{\rho}\right)\frac{1}{1-\vartheta L}\xi_{t+1} - \frac{1}{1-\rho}\left(1-\frac{\vartheta}{\rho}\right)\left(1-\frac{\hat{\lambda}}{\rho}\right)\frac{1}{1-\vartheta\hat{\lambda}}\frac{\rho+\vartheta-\rho\vartheta(L+\hat{\lambda})}{(1-\vartheta L)(1-\hat{\lambda}L)}\xi_t \\
&= \frac{1}{1-\rho}\left(1-\frac{\vartheta}{\rho}\right)\frac{1}{1-\vartheta L}\left(\varepsilon_{t+1} + \left(\rho - \left(1-\frac{\hat{\lambda}}{\rho}\right)\frac{1}{1-\vartheta\hat{\lambda}}\frac{\rho+\vartheta-\rho\vartheta(L+\hat{\lambda})}{1-\hat{\lambda}L}\right)\xi_t\right) \\
&= \frac{1}{1-\rho}\left(1-\frac{\vartheta}{\rho}\right)\left(\frac{\hat{\lambda}\vartheta(1-\rho\vartheta)}{\rho(\vartheta-\rho)(1-\hat{\lambda}\vartheta)}\frac{1}{1-\vartheta L} + \frac{\hat{\lambda}+\vartheta(\rho-\hat{\lambda})}{\rho(\vartheta-\rho)(1-\hat{\lambda}\vartheta)}\frac{1}{1-\hat{\lambda}L}\right)\varepsilon_{t+1}
\end{aligned}$$

Since $\hat{\lambda} < \rho$, the IRF of forecast error is always positive.

Proof of Proposition 7

If $\hat{\tau} = \tau = \infty$, then $\hat{\lambda} = \vartheta = 0$. As a result, $\omega_f = \frac{\hat{\rho}}{\rho}$, and $\omega_b = 0$. In this case, all agents receive the same signal, and there is no distinction between $\mathbb{E}_{i,t}[\cdot]$ and $\bar{\mathbb{E}}_t[\cdot]$. It follows that $K_{\text{CG}} = K_{\text{BGMS}}$.

To derive the K_{BGMS}, note that

$$\begin{aligned}
z_{t+1} - \mathbb{E}_{i,t}[z_{t+1}] &= \varepsilon_{t+1} + (\rho - \hat{\rho})z_t \\
\mathbb{E}_{i,t}[z_{t+1}] - \mathbb{E}_{i,t-1}[z_{t+1}] &= \hat{\rho}(z_t - \hat{\rho}z_{t-1})
\end{aligned}$$

It follows that

$$K_{\text{BGMS}} = \frac{\hat{\rho}(1-\rho\hat{\rho})(\rho-\hat{\rho})}{\hat{\rho}^2 + \hat{\rho}^4 - 2\rho\hat{\rho}^3}$$

Clearly, the sign of K_{BGMS} is the same as the sign of $\rho - \hat{\rho}$.

The law of motion of the forecast error is

$$z_{t+1} - \mathbb{E}_{i,t}[z_{t+1}] = \frac{1-\hat{\rho}\mathbb{L}}{1-\rho\mathbb{L}}\varepsilon_{t+1}.$$

The responses $\{\zeta_k\}_{k=1}^{\infty}$ are given by

$$\zeta_k = \rho^{k-1}(\rho - \hat{\rho}),$$

which are either all positive or all negative.

Proof of Proposition 8

We first consider the case with "higher-order doubts". The recursive formulation of individual consumer i's consumption choice is

$$c_{i,t} = \mathbb{E}_t[\xi_t] + \beta\mathbb{E}_t[c_{i,t+1}] + (1-\beta)\mathbb{E}_t[c_{t+1}]$$

As ξ_t is perfectly observed by consumer i, we guess the policy function is

$$c_{i,t} = a\xi_t,$$

for some constant a.

Under the assumption that agent i believes that other agents observe the fundamental shock with probability q, it follows that

$$\begin{aligned}
\mathbb{E}_{i,t}[c_{i,t+1}] &= \mathbb{E}_{i,t}[a\xi_{t+1}] = a\rho\xi_t \\
\mathbb{E}_{i,t}[\bar{\mathbb{E}}_t[\xi_t]] &= q\xi_t \\
\mathbb{E}_{i,t}[c_{t+1}] &= \mathbb{E}_{i,t}[\bar{\mathbb{E}}_t[a\xi_{t+1}]] = aq\rho\xi_t.
\end{aligned}$$

Substituting these expectations into consumers' optimal response leads to

$$a\xi_t = \xi_t + \beta a\rho\xi_t + (1-\beta)aq\rho\xi_t,$$

which further verifies our guess by setting the constant a as

$$a = \frac{1}{1-(\beta\rho+(1-\beta)q\rho)} < \frac{1}{1-\rho}.$$

In the economy without higher-order doubts but with mis-perceived $\hat{\rho}$, the aggregate outcome is

$$c_t = \frac{1}{1-\hat{\rho}}\xi_t.$$

The outcomes in the two economies are observationally equivalent iff

$$\frac{1}{1-\hat{\rho}} = \frac{1}{1-(\beta\rho+(1-\beta)q\rho)} \quad \rightarrow \quad \hat{\rho} = \rho - (1-\beta)\rho(1-q) < \rho$$

In terms of forecasts, in the economy with higher-order doubts,

$$\mathbb{E}_{i,t}[c_{t+1}] = \bar{\mathbb{E}}_t[c_{t+1}] = q\mathbb{E}^*_t[c_{t+1}].$$

where $\mathbb{E}^*_t[\cdot]$ is the perfect-information rational expectation operator.

Next, we consider the level-k thinking. The agents are assumed to observe the fundamental and to have the correct prior about its process but a mis-specified prior about the behavior of others: they are "level-k thinkers" for some finite integer $k \geq 0$. Level 0 agents are assumed to play $c_t = c^0_t \equiv 0$, for all t and for all ξ^t. Level 1 agents believe that other agents are level 0. They therefore play $c_t = c^1_t$, where c^1_t is given by the solution to

$$c^1_t = \xi_t + \beta\mathbb{E}_t[c^1_{t+1}]$$

Level 2 agents believe that other agents are level 1. They therefore choose $c_t = c^2_t$, where c^2_t is given by the solution to

$$c^2_t = \xi_t + \beta\mathbb{E}_t[c^2_{t+1}] + (1-\beta)\mathbb{E}_t[c^1_{t+1}].$$

Similarly, the aggregate outcome for level-k agent when $k > 0$ satisfies

$$c_t^k = \xi_t + \beta \mathbb{E}_t[c_{t+1}^k] + (1-\beta)\mathbb{E}_t[c_{t+1}^{k-1}].$$

We proceed by a guess-and-verify approach. Suppose that $c_t^k = a_k \xi_t$. Then for $k > 0$, a_k has the following recursive structure

$$a_k = 1 + \beta\rho a_k + (1-\beta)\rho a_{k-1}.$$

Using the fact $g_0 = 0$, we have for $k > 0$,

$$a_k = \frac{1}{1-\rho}\left(1 - \left(\frac{(1-\beta)\rho}{1-\beta\rho}\right)^k\right),$$

which has proved the conjecture.

Compared with the economy with mis-perceived $\hat{\rho}$, the aggregate outcomes are equivalent iff

$$\frac{1}{1-\hat{\rho}} = \frac{1}{1-\rho}\left(1 - \left(\frac{(1-\beta)\rho}{1-\beta\rho}\right)^k\right).$$

Since $\left(1 - \left(\frac{(1-\beta)\rho}{1-\beta\rho}\right)^k\right) < 1$, we have $\hat{\rho} < \rho$.

In terms of the forecast, in the level-k economy,

$$\mathbb{E}_{i,t}[c_{t+1}] = \overline{\mathbb{E}}_t[c_{t+1}] = a_{k-1}\rho\xi_t = \frac{a_{k-1}}{a_k}\mathbb{E}_t^*[c_{t+1}],$$

where $\frac{a_{k-1}}{a_k} < 1$.

Lastly, consider the cognitive discounting economy. We still proceed by a guess-and-verify approach. Suppose that the actual law of motion of c_t is

$$c_t = Rc_{t-1} + D\varepsilon_t,$$

and the perceived law of motion is

$$c_t = mRc_{t-1} + D\varepsilon_t.$$

Meanwhile, the perceived law of motion of ξ_t is

$$\xi_t = m\rho\xi_{t-} + \varepsilon_t.$$

Recall that the aggregate outcome is given by

$$c_t = \sum_{k=0}^{\infty} \beta^k \overline{\mathbb{E}}_t[\xi_{t+k}] + (1-\beta)\sum_{k=0}^{\infty}\beta^k \overline{\mathbb{E}}_t[c_{t+k+1}].$$

Using the mis-specified priors, we have

$$c_t = \frac{1}{1-\beta m\rho}\xi_t + (1-\beta)\frac{mR}{1-\beta mR}c_t,$$

which leads to the actual law of motion of c_t as

$$c_t = \rho c_{t-1} + \frac{1-\beta mR}{1-mR}\frac{1}{1-\beta m\rho}\varepsilon_t.$$

To be consistent with our guess, we have

$$R = \rho, \quad D = \frac{1}{1-m\rho}.$$

Compared with the economy with mis-perceived $\hat{\rho}$, the aggregate outcomes are equivalent iff

$$\frac{1}{1-\hat{\rho}} = \frac{1}{1-m\rho}, \quad \rightarrow \quad \hat{\rho} = m\rho < \rho.$$

In terms of the forecast, in the cognitive-discounting economy,

$$\mathbb{E}_{i,t}[c_{t+1}] = \overline{\mathbb{E}}_t[c_{t+1}] = m\mathbb{E}^*_t[c_{t+1}].$$

In all the three economies (higher-order doubts, level-k, cognitive discounting), the individual forecast is the same as the average forecast about the aggregate outcome , and it follows that $K_{\text{CG}} = K_{\text{BGMS}}$. In addition, in all the three economies,

$$c_t = \varphi\xi_t, \quad \text{and} \quad \overline{\mathbb{E}}_t[c_{t+1}] = \zeta\mathbb{E}^*_t[c_{t+1}] = \zeta\rho\varphi\xi_t,$$

for some constant φ and $\zeta \in (0,1)$. Therefore, we have

$$\text{Cov}(\text{Error}_t, \text{Revision}_t) = \text{Cov}(\varphi\xi_{t+1} - \zeta\rho\varphi\xi_t, \zeta\rho\varphi\xi_t - \zeta^2\rho^2\varphi\xi_{t-1}) = \varphi^2\rho^2\zeta(1-\zeta)\frac{1-\zeta\rho^2}{1-\rho^2},$$

which implies $K_{\text{CG}} = K_{\text{BGMS}} > 0$.

In addition, the law of motion of the forecast error is

$$\text{Error}_t = \varphi\frac{1-\zeta\rho\mathbb{L}}{1-\rho\mathbb{L}}\varepsilon_{t+1} = \varphi\left((1-\zeta)\frac{1}{1-\rho L} + \zeta\right)\varepsilon_{t+1},$$

and the corresponding IRF is always positive given $\zeta \in (0,1)$.