Title

Essays on Macroeconomic Expectations and Revised Data

Author

Paul, Tyler Louis

Publication Date

2024

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Essays on Macroeconomic Expectations and Revised Data

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Economics

by

Tyler Louis Paul

Committee in charge:

Professor Giacomo Rondina, Chair
Professor Juan Herreño
Professor Munseob Lee
Professor Valerie Ramey

2024

The Dissertation of Tyler Louis Paul is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2024

DEDICATION

In memory of my grandmama, Barbara Paul, and my cat, Nashi. I miss you both deeply.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

This dissertation would not have been possible without the enduring support from my wife. I truly appreciate all of the sacrifices you've made to enable my success. I owe gratitude to my parents for their constant encouragement and the many opportunities they created for me. I am thankful for the inspiration and joy my children have created while completing this work.

Thank you to my entire advisory committee for the wonderful guidance you've provided. I am especially grateful to Giacomo Rondina who was an excellent mentor both for my research and balancing the demands of graduate work with raising a young family. I am indebted to Valerie Ramey for her patient help throughout the years. I was fortunate to receive thoughtful and valuable advice from Munseob Lee and Juan Herreño.

I benefited from studying with Mitchell VanVuren, Pietro Spini, Alyssa Brown, Jacob Orchard, and Francesco Amodeo during graduate school. Their camaraderie and intellectual contributions have been invaluable to me.

Chapter 1 is currently being prepared for submission for publication; Tyler L. Paul. The dissertation author was the principal author on this paper.

Chapter 2 is currently being prepared for submission for publication; Tyler L. Paul. The dissertation author was the principal author on this paper.

Chapter 3 is currently being prepared for submission for publication; Tyler L. Paul. The dissertation author was the principal author on this paper.

VITA

2011	Bachelor of Science, Northeastern University
2011-2013	Analysis Group, Boston, MA
2014	Turning Point, Melbourne, VIC, Australia
2015–2016	Equifax, Melbourne, VIC, Australia
2016–2023	Teaching Assistant, University of California San Diego
2018	Research Assistant, University of California San Diego
2022	Instructor of Record, University of California San Diego
2023–2024	Assistant Professor, Teaching Stream, University of Toronto
2024	Doctor of Philosophy, University of California San Diego

FIELDS OF STUDY

Major Field: Economics (Specialization or Focused Studies)

Essays on Macroeconomic Expectations and Revised Data
Professor Giacomo Rondina

ABSTRACT OF THE DISSERTATION

Essays on Macroeconomic Expectations and Revised Data

by

Tyler Louis Paul

Doctor of Philosophy in Economics

University of California San Diego, 2024

Professor Giacomo Rondina, Chair

This dissertation studies the impact of data revisions on macroeconomic expectations. In the first chapter I show theoretically how revised data bias estimates of information frictions taken from surveys of expectations. The bias is due to the dynamic property of revised data which introduces additional past dependency into forecasts, breaking the informational symmetry between forecast errors and forecast updates. I propose a new estimation specification which corrects the bias and apply my method to data from the Survey of Professional Forecasters. I find the degree of information frictions is 33% different from existing estimates on average. Using my new estimates, I find equilibrium models of incomplete information more closely match the data.

In the second chapter, I study socially optimal data provision when the economy resembles a beauty-contest. The budget-constrained statistical office (S.O.) faces a trade-off: provide imprecise yet timely data about the present, or wait to provide more precise yet delayed data about the past. I find the S.O.'s welfare-maximizing choice of current versus revised signal precisions depends on the accuracy of private information and strength of strategic incentives. In some cases the S.O. offers a mixture of both signals, while in others it provides only one or neither signal in order to reduce the crowding-out of private information caused by the coordination incentive. The S.O. treats the revised signal as an inferior good, reducing its precision in favor of the current signal when the budget limit rises.

In the third chapter, I evaluate the interaction between revised data and higher-order dynamics when estimating information frictions. I prove analytically that the existing estimation method is valid when the fundamental possesses additional past dependency relative to the case in chapter one. I find additional bias terms in estimates of information frictions in the presence of revised data and higher-order dynamics. I propose a correction which eliminates the bias and does not depend on the specific time series properties of the data. I use simulation to study the nature of this new bias, finding it to be quantitatively small relative to the bias caused solely by revised data.

Chapter 1

Measuring Information Frictions with Surveys of Expectations: How Revised Data Bias Estimates

I study how data signals of the past affect the prevailing method of estimating information frictions. Existing literature measures the relationship between average forecast errors and forecast updates to estimate the degree of information frictions, or slow response of forecasts to news. I introduce a revised data signal into the noisy information model and find this theoretic relationship now suffers from omitted variable bias and cannot be estimated by OLS. The bias is not due to noise in the signal but due to the dynamic property of the revised data which introduces additional past dependency into forecasts and breaks the Markovian flow of information. I use the model to propose a new specification which corrects the bias. When I estimate the corrected specification on data from the Survey of Professional Forecasters, I find the degree of information frictions is 33% different from existing estimates on average. I apply my results toward two applications to demonstrate the economic significance of this bias. Under my corrected estimates I find incomplete information theories more closely match the data, while the efficacy of forward guidance as a monetary policy tool is significantly reduced.

1.1 Introduction

Most aggregate data series in macroeconomics are imperfectly estimated using surveys and statistical sampling. These estimates can change overtime as more information becomes available or new classification schemes are used, which generates revised estimates of existing published values.[1] A large literature has studied the value of using revised data to improve real-time forecasts, repeatedly finding increased performance relative to models which ignore revisions (Clements and Galvão, 2017). This paper applies these insights towards a question at the heart of modern macroeconomics: how are expectations formed?

Canonical workhouse models in macro assume full information and rational expectations (FIRE): agents freely observe all relevant variables and form their forecasts optimally. Using surveys of forecasters past researchers have empirically tested these assumptions, generally finding evidence against either or both components of FIRE (Pesaran and Weale, 2006). Recent work has followed the contribution of Coibion and Gorodnichenko (2015), henceforth CG, who develop a test of the full information assumption based on predictability of the forecast error. When the FIRE assumption holds, forecast errors should be unpredictable; otherwise the forecast was not optimally using all available information and could be improved. CG show how models with either sticky or noisy information predict that forecast errors averaged across all agents are predictable by average forecast updates, which are cross-period revisions to a forecast. The degree of predictability maps directly to a measure of information frictions, or under-reaction to informative signals. When they estimate their test using OLS they find just under half of the information available to forecasters is used in their forecasts.

The CG test is a seminal tool for capturing information frictions directly from data. However, it is derived from an information set which does not allow signals of the past. Given the usefulness of revised data in forecasting, this paper asks: how are estimates of information frictions affected by revised data? I make two key contributions. First, I demonstrate theoretically how including a revised signal in the noisy information framework biases the CG test and renders regression estimates unable to identify the degree of information frictions. The bias can be economically large, depending on parameters, and has a complex interaction with relative signal

[1] These revisions can be large: over the past two decades, the first estimate of the monthly change in non farm payroll employment had a 90% confidence interval of $\pm 115,000$ jobs; the average level of this estimate was 128,000 (Mester, 2016). Another example is the well-known revision during the Great Recession, where the 2008Q4 real GDP growth was initially estimated at -3.8%, revised one quarter later to -6.3%, and finally lowered to -8.4% where it stays today.

precisions and persistence. Furthermore, the structural mapping between regression estimates and information frictions is altered due to revised data, introducing additional error in the recovered size of the friction. Second, I propose a new specification of the test which corrects the bias and restores identification. I estimate my specification and compare with estimates from CG to show the empirical relevance of the bias. I find the bias ranges from -0.49 to 0.28 percentage points across a variety of macroeconomic variables and forecasting horizons.

To understand the source of the bias, first consider how forecasts are updated. Upon receiving a new contemporaneous signal, the forecaster learns about the current period's innovation for the first time, while also shedding light on past innovations. Optimal forecasting behavior implies the current period forecast will include some portion of the new signal, according to its accuracy, plus a fraction of the prior forecast, weighted by the signal accuracy and persistence. In effect, the flow of information follows the process of the state: past information is carried forward similarly to how the state evolves through time, but down-weighted by the trustworthiness of new information. The dependent variable in the CG test, forecast error, also inherits this same Markovian flow from the state.

When forecasters also observe a revised signal, the flow of information is broken since forecasts contain extra learning about the past. The forecast update thus includes an additional 'bump' of past dependence, which changes the weights on all previous innovations. This does not affect the forecast error, though, since the error compares the entire history of forecasts with the realization, not the change across periods as for the forecast update. This means the new past dependence ends up in the regression residual as an omitted variable and biases estimates of the relationship between forecast errors and forecast updates.

I use the noisy information model with the expanded information set to identify an additional regressor to correct the bias in the CG test. Alongside forecast updates I include the revised signal innovation, or the difference between the revised signal and the lagged expectation of it. This term captures the extra past dependence of forecasts (or the 'bump' in the flow of information) and removes both the lagged fundamental and its noise from the error term. The correction term also restores the mapping between the estimated coefficients and the degree of information frictions.

In order to test my correction term I require a revised data signal. Typically, a data revision occurs when additional data become available at a later date and are used to provide a more accurate estimate of the fundamental. These revisions are common for data series in the national income and product accounts (NIPA), since most series are constructed from surveys of private entities which can have large variation in response time. The most recent estimates are assumed to be the most accurate, though my findings do not require this assumption. In terms of the news versus noise literature of Mankiw, Runkle and Shapiro (1984), Mankiw and Shapiro (1986), and Aruoba (2008), these revisions are forecastable given contemporary information and thus considered noise.

There is an additional type of data revision that I study empirically due to the structure of the expectations data. The publicly available Survey of Professional Forecasters (SPF) is conducted quarterly by the Philadelphia Federal Reserve. Responses are sent in the middle of the quarter, while forecasts are made on a quarterly basis. This means that for some variables which are released monthly, like the unemployment rate, forecasts are based on incomplete data for the current quarter (namely, the final month is missing). This also applies to high-frequency financial variables, like the interest rate on Treasury bills, which is known with certainty each day. I consider the initial signal of these variables to be what was known by forecasters at the time they returned the survey, while the revised signal is the actual quarterly value known at the end of the quarter. In this way, I may test the model with variables which are measured precisely and thus do not have the typical revision process.[2] In total I will test 14 aggregate data series for the influence of revised data: six from the NIPA and eight variables with a greater frequency than quarterly.

When taking the model to the data, I have two goals: to assess the theory's validity and to

[2] These revisions would be classified as news: the revision is correlated with the true value, but unpredictable given contemporary information.

re-estimate information frictions. To validate the model I test if the coefficient on the correction term is statistically different from zero, even after conditioning on forecast updates, which implies forecasters pay attention to revised data. This would be contrary to the investigations of CG and others who find no other variable helps predict forecast errors when forecast updates are included.[3] In my benchmark specification, I find a non-zero coefficient estimate for 12 out of 14 variables, which is strong support of the theory. Across other specifications with different forecasting horizons the average number of variables with a non-zero coefficient drops to 8 out of 14. The NIPA variables in particular are most likely to have a correction term coefficient estimate of zero as the horizon changes, unlike the higher-frequency variables which nearly always are non-zero. These results indicate forecasters use revised data as predicted by the model, though revised data may be more important for non-NIPA variables.

Next I re-estimate the degree of information frictions and statistically compare my estimate to one generated by the CG specification. I find six variables with a different estimate of frictions between the two models under my benchmark specification. The difference amongst these variables is nearly a third on average in both directions: CPI inflation, housing starts and the unemployment rate are underestimated by the CG model by 20 percentage points (a 28% difference), while two financial variables are overestimated by 35 percentage points (a 43% difference). As an example, existing estimates suggest 80% of new CPI inflation data are incorporated into forecasts while my corrected estimates say only 57% of these data are used. The difference of 23 percentage points corresponds to a 29% bias in the estimate of information frictions.

Interestingly, all of these six variables are non-NIPA, a pattern that generally holds across forecasting horizons. Combined with the findings from the model validation exercise it appears there is a special interaction between NIPA data and data revisions. Alternatively, it may be the 'constructed' nature of non-NIPA data revisions that is unrealistically influencing the estimates. I provide two possible explanations. First, the natural revisions of non-NIPA data

[3] One specification in CG found the unemployment rate significant, but the cause is unknown.

can be considered a closer match to the theory, since the mechanism depends on how the flow of information is altered by the revised signal. Non-NIPA revisions are highly informative to forecasters, so one would expect them to break the flow of information and thus bias existing estimates of information frictions. Second, NIPA revisions occur out of sync with the forecaster survey so that the time between when forecasters observe the first estimate and most recent revision is several months. During these months multiple revisions occur, so my ability to match changes in information with changes in expectations is limited. In summation, my empirical results provide strong evidence to believe revised data are an important feature of expectations and estimating information frictions, though their impact may be weakened on NIPA data due to their low frequency.

As an extension of the theory I apply my correction term toward the salient question of individual rationality recently raised by Kohlhas and Broer (2019) and Bordalo et al. (2020). These authors find evidence of forecast error predictability at the individual level, violating the rational expectations assumption. While the bias due to revised data could explain these findings in theory, estimating my specification with the correction term at the individual level instead corroborates them. Indeed, the degree of overreaction at the individual level is typically higher once correcting for revised data.

To demonstrate the importance of accurate estimates I study two applications and compare results from my estimates with CG. In the first I assess the performance of a modern New Keynesian model with incomplete information from Angeletos and Huo (2021). The goal of introducing information frictions is to better match the persistence observed in data estimates of impulse response functions. When I plot the IRFs from the model I find the curve is significantly more hump-shaped when using my estimate of information frictions relative to CG. Evaluation of the model's performance with only CG's estimate may lead one to reject the power of information frictions, which would be an incorrect conclusion if my estimate was closer to the truth. In the second application I evaluate a stylistic model of forward guidance courtesy of Angeletos and Lian (2018). When I recreate their results I find the consumption response to a forward announcement of interest rates is sharply reduced under my estimate, more than five times lower than the response in CG. Using the most accurate degree of information frictions thus matters for determining the usefulness of unconventional monetary policy.

This paper contributes to three strands of literature. The first contains papers which empirically test FIRE and recover estimates of information frictions. Many authors have tested FIRE since Muth (1960) put forward model-consistent expectations as a standard (Pesaran and Weale, 2006), but the recent methodological contributions from Coibion and Gorodnichenko (2012, 2015) sparked a flurry of research using survey data. Kumar et al. (2015) and Coibion, Gorodnichenko and Kumar (2018) study how firms in New Zealand form their expectations, particularly in response to central bank activity. Coibion et al. (2021*b*) designed a survey eliciting higher-order expectations, which they use to test alternative theories of incomplete information like level-*k* thinking or heterogeneous priors. Cutting edge studies create randomized controlled trials to test the causal impact of various information treatments on firm and individual expectations and behaviors (Coibion, Gorodnichenko and Ropele (2020), Coibion et al. (2021*a*), Coibion, Gorodnichenko and Weber (2022)). There is an abundance of new survey data and additional tests of expectations that could be affected by revised data similar to the original CG specification.

Within this set of literature, my paper fits most directly amongst a recent wave of papers that examine the validity of the CG test. Bordalo et al. (2020) apply the CG test to individual forecasters and find evidence of individual overreaction to shocks. The authors take this as evidence against rational expectations and suggest a non-Bayesian model, similar to Kohlhas and Broer (2019). In theory, the bias I identify due to revised data could explain over-reaction at the individual level, however when I estimate my corrected specification on individual forecasters I find even stronger evidence for over-reaction. Goldstein (2019) derives an alternative approach to measuring information frictions that is feasible at the individual level by exploiting forecast dispersion about the mean. His information framework does not include revised data, which may similarly affect the test's validity. Gemmi and Valchev (2022) find bias in surveys of forecasters subject to strategic considerations and common noise. My paper does not include strategic coordination, but I find a bias nonetheless which is not due to noise but due to the dynamic property of revised data. Goldstein and Gorodnichenko (2022) add a noisy future signal to each forecaster's information set. They show how forward information changes the persistence of forecasts relative to the underlying state, a similar result to what I find with revised information. Their paper, however, does not consider how forward information may bias estimates of information frictions.

My contribution also speaks to the literature which analyzes data revisions and builds forecasting models with real-time data. Concern over the vintage used in forecasting models extends back to Cole (1969) who found that using revised estimates of GNP produces more accurate forecasts than using preliminary estimates. Howrey (1978, 1984) built upon this insight by providing a framework for incorporating preliminary and revised data together for optimal economic forecasting. The modern success of this literature started with the efforts of Croushore and Stark (2001) who developed a real-time data set of macroeconomic variables. The data are arranged into vintages according to the latest known observation for each period of a time series, enabling econometricians to quickly construct information sets as they were available at a given point in term. In Stark and Croushore (2002) the authors compare ARIMA forecasting models which use the latest vintage versus the real-time vintage and find large differences in average errors. Since then several authors have created forecasting models which incorporate multiple vintages. Jacobs and van Norden (2011) construct a state-space model which separately allows for news, noise and spillover effects of data revisions. Kishor and Koenig (2009) generalize Howrey (1978) in a VAR framework where the data are filtered by the statistical agency before publication, as in Sargent (1989). Galvão (2017) jointly estimates the parameters of a DSGE model with the data revision process, improving the forecast accuracy of the DSGE model. While these papers use sophisticated modeling techniques to harness the value of revised data, I study a simple model of data revisions and take it as sufficient that incorporating one additional vintage of data improves forecast performance. To my knowledge, this is the first paper to apply the forecasting benefits of revised data towards questions of incomplete information.

By providing new estimates of information frictions my paper contributes to the theoretical literature studying incomplete information models. An early example is Lucas (1972), which described the economy as a series of islands that hold private information. More recent approaches follow either the sticky information theory of Mankiw and Reis (2002), rational inattention of Sims (2003), or the noisy information framework of Woodford (2001). Since my paper uses noisy information my results are most applicable to literature in this vein. Nimark (2014) considers an information framework with multiple sources of private and public noise, and shows how they can produce realistic business cycle behavior. Melosi (2016) proposes a DSGE model where firms receive noisy private signals of aggregates, which ultimately makes inflation expectations respond more slowly than inflation itself. While these models are estimated directly on survey data, they could in theory be calibrated by my estimates of information frictions. More recent advancements from Angeletos and La'O (2020), Huo and Pedroni (2020), and Angeletos and Lian (2022) recast the incomplete information model into a general global games framework which sharpens the connection between estimates of information frictions and their aggregate implications. For this reason I chose two recent papers by Angeletos and Huo (2021) and Angeletos and Lian (2018) to demonstrate the significance of my findings.

The paper proceeds as follows. Section 1.2 presents the model and analyzes the bias caused by a revised signal. Intuition is developed initially through the special-case of a perfectly-revealing revised signal Section 1.3 describes the expectations and macroeconomic variable data, with special attention paid to the revision process. Empirical results are in section 1.4, which I use in two applications in section 1.5.

1.2 Model and Analytic Characterization of Bias

I study a noisy information model similar to CG and Woodford (2001), but adapted to include a revised data signal. The fundamental x_t follows an AR(1) process with serially uncorrelated errors. There is a continuum of forecasters indexed by i who wish to predict the fundamental. Each forecaster privately receives a current signal of the fundamental based on their own knowledge and experience.[4] New to their model, I add an additional revised signal that all forecasters observe, r_t, which is informative of the previous period's fundamental:

$$\text{State: } x_t = \rho x_{t-1} + w_t$$

$$\text{Current signal: } y_{it} = x_t + \psi_{it}$$

$$\text{Revised signal: } r_t = x_{t-1} + \nu_t$$

where $0 < \rho < 1$, $w_t \sim N(0, \tau_w^{-1})$, $\psi_{it} \sim N(0, \tau_\psi^{-1})$, and $\nu_t \sim N(0, \tau_\nu^{-1})$. The idiosyncratic term ψ_{it} is uncorrelated across both agents and time, and represents the influence of the forecaster's local experiences and expertise as in Berger, Ehrmann and Fratzscher (2011). The common noise ν_t represents the imprecision inherent in revised data publications from the statistical authority. Importantly I allow for a perfectly-revealing revised signal, when $\tau_\psi \to \infty$, which provides useful intuition and helps match some data series. The common noise term has no serial correlation. Finally, computing the revision as $y_{it-1} - r_t = \psi_{it-1} - \nu_t$ which is pure noise, one potential interpretation of data revisions. I show in appendix A.1.2 that a model of data revisions as news leads to qualitatively similar results as I find with the noise specification.

Each forecaster forms their best prediction of the fundamental conditional on her information set, Ω_{it}, which consists of all current and revised signals revealed through time t: $\Omega_{it} = \{y_{it}, y_{it-1}, \ldots, y_0, r_t, r_{t-1}, \ldots, r_1\}$. For notational convenience I define $\mathbb{E}_{it}[x_{t+h}] = \mathbb{E}[x_{t+h}|\Omega_{it}]$. Since all innovations are assumed to be normal, I can use the Kalman filter to determine the optimal forecasting rule. I show in appendix A.1.1 how to derive the following forecasting

[4] CG study a current signal with both common and idiosyncratic noise in their appendix. They find the common noise term biases their estimate of information frictions downward, making their results an underestimate of the degree of information frictions. I will show that common noise in my model behaves similarly, but is not crucial to my findings.

formula:

$$\mathbb{E}_{it}[x_{t+h}] = \mathbb{E}_{it-1}[x_{t+h}] + \rho^h K(y_{it} - \mathbb{E}_{it-1}[y_{it}]) + \rho^h K_r(r_t - \mathbb{E}_{it-1}[r_t]) \tag{1.1}$$

The parameters K and K_r are the steady-state Kalman gains, or the weights placed on the current and revised signals when forecasting the current fundamental x_t.[5] CG use their forecasting formula to derive a relationship between average forecast errors, $x_{t+h} - \mathbb{E}_t[x_{t+h}]$, and average forecast updates, $\mathbb{E}_t[x_{t+h}] - \mathbb{E}_{t-1}[x_{t+h}]$. The forecast update measures how the forecast of x_{t+h} changes between periods $t-1$ and t (that is, after observing new signals in period t). Intuitively, this captures the relevance of new information learned in period t. If there were no signals observed in period t, then the forecast update should be 0; today's forecast of x_{t+h} would be the same as yesterday's forecast. To find the relationship between forecast errors and revisions, I will manipulate the optimal forecasting rule, equation (1.1).

1.2.1 Case 1: Irrelevant revised signal

To match the setting used by CG I initially assume the revised signal is not useful for forecasting because it is infinitely noisy, or $\tau_v \to 0$. In this case the weight on the revised signal $K_r = 0$, so the optimal forecast becomes:

$$\mathbb{E}_{it}[x_{t+h}] = \mathbb{E}_{it-1}[x_{t+h}] + \rho^h K(y_{it} - \mathbb{E}_{it-1}[y_{it}])$$

From this equation I plug in y_{it}, average across forecasters to eliminate idiosyncratic noise, subtract $K\mathbb{E}_t[x_{t+h}]$ from both sides and rearrange to get equation (9) from CG:

$$x_{t+h} - \mathbb{E}_t[x_{t+h}] = \left(\frac{1-K}{K}\right)(\mathbb{E}_t[x_{t+h}] - \mathbb{E}_{t-1}[x_{t+h}]) + \gamma_{t+h} \tag{1.2}$$

[5]Because the fundamental is stationary, the Kalman gain matrix will converge. See Hamilton (1994).

where $\gamma_{t+h} = \sum_{j=1}^{h} \rho^{h-j} w_{t+j}$ is a linear combination of future fundamental shocks. Equation (1.2) relates average forecast errors to average forecast updates. Averaging across forecasters eliminates the idiosyncratic noise term ψ_{it} in y_{it}.[6] Since γ_{t+h} consists of shocks starting in $t+1$, it is not known by forecasters in period t, so it must be uncorrelated with the forecast update.[7]

From this structural relationship, CG suggest the following econometric specification:

$$x_{t+h} - \mathbb{E}_t[x_{t+h}] = \beta\left(\mathbb{E}_t[x_{t+h}] - \mathbb{E}_{t-1}[x_{t+h}]\right) + \varepsilon_t \tag{1.3}$$

where ε_t is the regression residual. Comparing this with (1.2) shows $\varepsilon_t = \gamma_{t+1}$, a series of shocks occurring after period t and thus uncorrelated with the regressor. Therefore (1.3) can be estimated consistently with OLS; estimates of β recover the structural coefficient in equation (1.2), which I will call $\bar{\beta}_{CG} = \frac{1-K}{K}$.[8] Through this connection, one can get an estimate of K, the degree of information frictions: $\hat{\beta} \to \bar{\beta}_{CG}$, or $\hat{K} = \frac{1}{1+\hat{\beta}}$.

The reason this relationship holds is the following logic. The forecast update represents all new information learned in period t relative to $t-1$. However, that information is imprecisely observed due to the noise in signal y_{it}, so forecasters will only partially incorporate it into their new predictions by down-weighing the signal by $K < 1$, with the rest of the weight $(1-K)$ placed on their prior belief. The forecast update thus contains only a fraction of the new information; it is this less-than-full updating which creates the predictability of forecast errors by past forecasts.

To foreshadow my results, I will expand the key terms in (1.2) into two components: information about period t and information from periods $t-1$ and earlier.

Proposition 1 (Forecast error and revision composition) *When the revised signal is irrelevant, the average forecast error and average forecast update can be represented with the*

[6] Any current period common noise term would enter the right-hand side negatively.

[7] This term disappears when $h = 0$.

[8] Since any common noise term would enter ε_t negatively, estimates of $\hat{\beta}$ would be biased downwards and thus underestimate the degree of information frictions. CG argue that any finding of $\hat{\beta} > 0$ is strong evidence of information frictions when common noise is present.

following recursive equations:

$$x_{t+h} - \mathbb{E}_t[x_{t+h}] = \rho^h(1-K)w_t + \rho^{h+1}(1-K)(x_{t-1} - \mathbb{E}_{t-1}[x_{t-1}]) + \gamma_{t+h}$$
$$\mathbb{E}_t[x_{t+h}] - \mathbb{E}_{t-1}[x_{t+h}] = \rho^h K w_t \qquad + \rho^{h+1}K(x_{t-1} - \mathbb{E}_{t-1}[x_{t-1}])$$

All proofs are relegated to appendix A.1.5. Both equations are linear combinations of the current shock w_t and last period's forecast error, with the respective weights depending on the degree of signal precision. The forecast update places a weight of K on the current shock since w_t is noisily observed from y_{it}, so only some fraction K is incorporated into period t forecasts. This implies that a fraction $(1-K)$ of the current shock is missing from the period t forecast, and thus shows up in the forecast error. The second term contains last period's learning from y_{it-1} and thus the rational errors made predicting shocks from $t-1$ and earlier.[9] These errors are discounted according to how far back they are so that very old forecast errors are less important for today's error and forecast update. The key insight from proposition 1 is that the underlying process for both forecast error and forecast update is the same in terms of fundamental shocks. It is the difference in weights ascribed to each shock due to the noisy signal that makes the relationship predictable.

CG stress that this predictability only arises after aggregating across forecasters; otherwise, individuals could exploit the relationship to reduce their own forecast error, implying their original forecast was not optimal. Recent empirical work, such as Ryngaert (2017) and Bordalo et al. (2020), has tested this relationship on an individual level and found predictability, casting doubt on the rational expectations framework. I will show that my results could address this discrepancy in theory, however individual predictability is still empirically present after correcting for revised data.

Next I will allow $\tau_v \in (0, \infty)$ and show how the regression residual is no longer orthogonal to the forecast update.

[9] Indeed the forecast error is a difference equation which I will solve backwards later.

1.2.2 Case 2: General revised signal

Let the revised signal have forecasting value yet be observed with noise so the optimal forecasting equation is (1.1), reproduced below:

$$\mathbb{E}_{it}[x_{t+h}] = \mathbb{E}_{it-1}[x_{t+h}] + \rho^h K(y_{it} - \mathbb{E}_{it-1}[y_{it}]) + \rho^h K_r(r_t - \mathbb{E}_{it-1}[r_t])$$

The goal is to recreate relationship (1.2) under this new information framework and characterize the bias. Starting from the optimal forecasting rule (1.1), I perform similar manipulations as for case 1. However, now there is a new term that doesn't match (1.2): the revised signal innovation $(r_t - \mathbb{E}_{t-1}[r_t])$. Under the CG relationship, this term will be captured by the forecast update in an identical manner to the current signal innovation $(y_{it} - \mathbb{E}_{t-1}[y_{it}])$. I will show that the revised signal innovation is fundamentally different, and thus representing it by the forecast update requires an additional term to appear. First expand this signal innovation and relate both innovations as follows:

$$\rho(r_t - \mathbb{E}_{t-1}[x_{t-1}]) = (y_{it} - \mathbb{E}_{t-1}[x_t]) - w_t + (\rho \nu_t - \psi_{it})$$

The above equation states that after accounting for the degree of persistence, the difference in information between the two signal innovations is w_t, the fundamental innovation in the current period, plus the noise terms from both signals. All of the other information embedded in y_{it} and r_t is identical. In order to recreate (1.2) with a revised signal I must summarize the mutual information content across both innovations, but to do that requires an additional term in w_t to account for the difference in information. The following proposition uses the above result to find the structural relationship in the presence of revised data.

Proposition 2 (Structural equation with revised data) *When the revised signal is informative,* *the true structural relationship between average forecast errors and average forecast updates is:*

$$x_{t+h} - \mathbb{E}_t[x_{t+h}] = \left(\frac{\rho(1-K)-K_r}{\rho K + K_r}\right)(\mathbb{E}_t[x_{t+h}] - \mathbb{E}_{t-1}[x_{t+h}]) + \left(\frac{\rho^h K_r}{\rho K + K_r}\right) w_t + \gamma_{t+h} - \left(\frac{\rho^{h+1} K_r \nu_t}{\rho K + K_r}\right) \tag{1.4}$$

There are two key differences when relating (1.4) to the structural equation (1.2). First, the coefficient on forecast updates (which I will call $\bar{\beta}_{REV}$) is now a combination of K, K_r, and ρ. Second, the structural residual contains a term in ν_t, the common noise of the revised signal, and a term in w_t, the current period fundamental innovation. In terms of the econometric specification (1.3), the first difference means estimates of $\hat{\beta}$ no longer directly map into an estimate of information frictions, K, since estimates of K_r and ρ are needed as well. The second difference implies there is correlation between the regressor and residual through w_t, so the OLS assumption of orthogonality no longer holds. Next I will characterize the bias when estimating (1.3) with OLS.

1.2.3 Bias Characterization

If one were to estimate equation (1.3) when revised data are present, then the true structural β one is attempting to recover would be $\bar{\beta}_{REV} = \frac{\rho(1-K)-K_r}{\rho K+K_r}$. However, the presence of w_t in the error term ε_t means OLS estimates of β will be biased due to correlation with average forecast updates. Running OLS on (1.3) when (1.4) is the correct model will give the following estimate of β:

$$\hat{\beta} = \beta + \frac{\mathbb{E}\left[(\mathbb{E}_t[x_{t+h}] - \mathbb{E}_{t-1}[x_{t+h}])\left(\left(\frac{\rho^h K_r}{\rho K+K_r}\right) w_t + \gamma_{t+h} - \left(\frac{\rho^{h+1} K_r \nu_t}{\rho K+K_r}\right)\right)\right]}{\mathbb{E}[\mathbb{E}_t[x_{t+h}] - \mathbb{E}_{t-1}[x_{t+h}]]^2}$$

$\mathbb{E}[(\mathbb{E}_{it}[x_{t+h}] - \mathbb{E}_{it-1}[x_{t+h}])\gamma_{t+h}] = 0$ since γ_{t+h} only contains shocks from the future. That leaves two terms with opposite sign: a positive term in w_t and a negative term in ν_t. The following proposition analytically calculates the size of the bias.

Proposition 3 (Bias equation) *If the true information structure contains a revised public signal (e.g. $\tau_v \in (0,\infty)$) then estimating specification (1.3) with OLS will produce a biased estimate of β. The size of the bias is:*

$$Bias = \frac{(1-G^2)\left(\frac{KK_r\tau_w^{-1}-\rho K_r^2\tau_v^{-1}}{\rho K+K_r}\right)}{\tau_w^{-1}\left(K^2+2K_rKG+K_r^2\right)+K_r^2\tau_v^{-1}\left(1-2G\rho+\rho^2\right)} \tag{1.5}$$

where $G=(\rho-\rho K-K_r)$.

The direction of the bias is ambiguous; if $\frac{\tau_v}{\tau_w} > \rho\frac{K_r}{K}$ the bias will be positive and $\hat{\beta}$ will be an over-estimate. Given the true structural coefficient $\bar{\beta}_{REV}$ I can compute the relation between $\hat{\beta}$ and the implied estimate of K:

$$K = \frac{1}{1+\hat{\beta}} - \frac{K_r}{\rho} \tag{1.6}$$

Equation (1.6) collapses to $\frac{1}{1+\beta}$ when $K_r = 0$, as in the CG case with an irrelevant revised signal. Larger estimates of $\hat{\beta}$ lead to smaller implied K, so a positive bias makes estimating (1.3) more likely to find higher degrees of information friction than there truthfully are.

The size of the bias is affected by the relative precisions of the fundamental, the noise terms, and the persistence, but importantly not the forecast horizon. Figure 1.1 plots equation (1.5) for varying levels of τ_v and ρ while holding the current signal precision ratio $\frac{\tau_\psi}{\tau_w}$ at a moderately high level.[10] The graph shows that the bias is largest for low (but non-zero) degrees of persistence, a pattern which is preserved as $\frac{\tau_\psi}{\tau_w}$ changes. This result seems counter-intuitive: if the bias is only present due to the revised signal (i.e. $K_r \neq 0$), then decreasing the importance of the revised signal by lowering ρ should decrease the bias. To better understand why this happens, I study a simplified model where the revised signal is perfectly-revealing: $\tau_v \to \infty$.[11]

[10] As the current signal precision ratio rises, both the bias and the difference due to persistence decrease, since the revised signal becomes relatively less important.

[11] This simplification is with some loss of generality, which I will make clear in the following section. The basic intuition is the same in the general case.

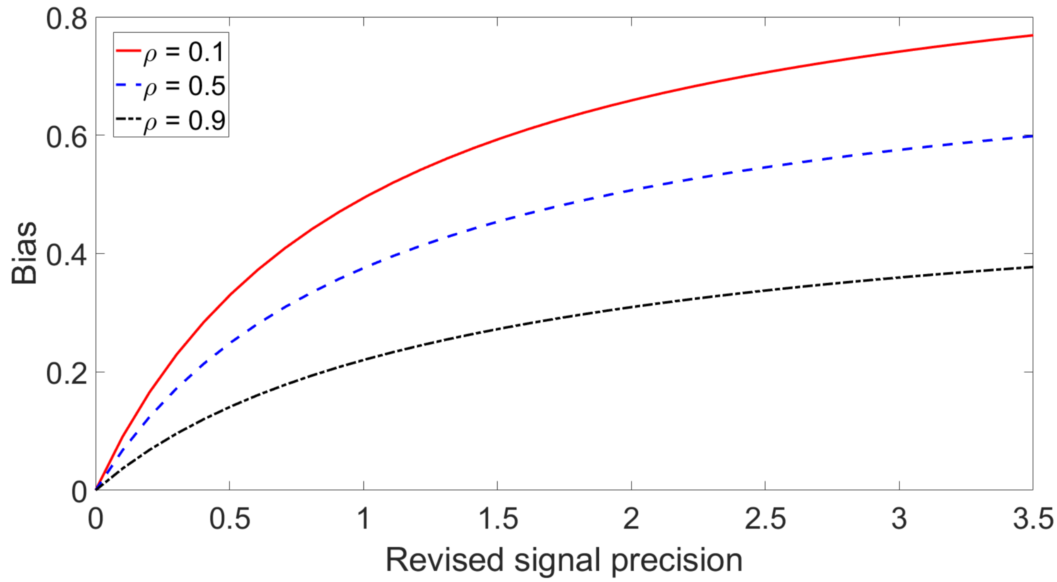

Figure 1.1. Bias when current signal precision is high. This graph plots the bias from equation (1.5) as the degree of revised signal precision changes for different levels of persistence ρ. The degree of current signal precision $\frac{\tau_\psi}{\tau_w} = 1$. The bias is largest when persistence is low.

Intuition: perfectly-revealing revised signal

A perfectly-revealing revised signal greatly simplifies the forecasting equations which allows us to neatly identify the source of the bias. Additionally, many authors explicitly or implicitly assume that the state is perfectly-revealed with a lag. For instance, the theoretical literature often uses a perfectly-revealing revised signal to limit the extent to which individuals have differentiated information sets, thus permitting a solution to incomplete information models (Lucas, 1972). More relevantly, most empirical papers implicitly makes this assumption by choosing a specific data release as the 'final' version, or empirical counterpart to x_t. Studying a perfectly revealing revised signal therefore helps connect this paper's contribution to the wider literature.

When x_{t-1} is known in period t, the problem becomes a static forecasting exercise of predicting w_t.[12] The Kalman gain parameters simplify to $K = \frac{\tau_\psi}{\tau_\psi + \tau_w}$, and $K_r = \frac{\rho \tau_w}{\tau_\psi + \tau_w}$, which

[12] I could proceed by forecasting x_t using the modified signal $\tilde{y}_{it} = y_{it} - \rho r_t$, but I prefer to keep the signals intact to demonstrate the flow of information. The resulting conclusions are the same.

allows recomputing of the individual forecast (1.1):

$$\mathbb{E}_{it}[x_t] = K(\rho x_{t-1} + w_t) + K_r x_{t-1} + \xi_{it} = \rho x_{t-1} + K w_t + \xi_{it} \tag{1.7}$$

where $\xi_{it} = K\psi_{it}$. Forecasts no longer depend on the prior belief $\mathbb{E}_{it-1}[x_t]$ since r_t perfectly reveals the state from the previous period. Despite this, the weight on r_t (i.e. K_r) is less than 1. This is because the current signal y_{it} also contains information about x_{t-1}, but since that signal is noisy agents only learn a fraction $\rho K x_{t-1}$ from y_{it}. Efficiently using r_t alongside y_{it} demands $K_r = \rho(1-K)$ because agents perfectly know x_{t-1}.

Now consider the process of forecast updating. The difference in information content between today's nowcast and yesterday's forecast should reflect the new information acquired today. If I expand (1.7) to show the information in each signal, I can keep track of the source of new information in the forecast update:

$$\begin{aligned}
\mathbb{E}_{it}[x_t] - \mathbb{E}_{it-1}[x_t] &= K(\underbrace{\rho^2 x_{t-2} + \rho w_{t-1}}_{\text{Info on } x_{t-1} \text{ from } y_{it}} + w_t) + K_r(\underbrace{\rho x_{t-2}}_{\text{Info on } x_{t-2} \text{ from } r_t} + w_{t-1}) \\
&\quad - \rho(K(\underbrace{\rho x_{t-2} + w_{t-1}}_{\text{Info on } x_{t-1} \text{ from } y_{it-1}}) + \underbrace{K_r x_{t-2}}_{\text{Info on } x_{t-2} \text{ from } r_{t-1}}) + \tilde{\xi}_{it} \\
&= K w_t + K_r w_{t-1} + \tilde{\xi}_{it}
\end{aligned} \tag{1.8}$$

where $\tilde{\xi}_{it} = \xi_{it} - \rho\xi_{it-1}$. Today's nowcast depends on y_{it}, which contains the current fundamental shock w_t as well as all past fundamental shocks in x_{t-1} due to the AR(1) nature of the state. However, the forecaster already learned of these past shocks in the previous period from y_{it-1}, so updating the forecast removes this information. The same happens with common information about x_{t-2} in r_t and r_{t-1}. All that remains is w_t as learned from y_{it} and w_{t-1} as learned from r_t.[13]

[13] The form of (1.8) is especially simple due to the perfectly-revealing revised signal. In the general case when $\tau_v < \infty$, the forecast update will still depend on all past shocks since observing y_{it} allows forecasters to update their predictions of the past. When $r_t = x_{t-1}$, there is no guesswork involved in the past beyond w_{t-1}, so we are left with (1.8). See section 1.2.3 for the general case.

The presence of w_{t-1} in the forecast update is due to forecasters using the revised signals to perfectly infer all past shocks. The forecaster knows x_{t-1} from r_t and x_{t-2} from r_{t-1}. Combining both signals is akin to updating her belief of the past shocks: $r_t - \rho r_{t-1} = w_{t-1}$. The forecast update therefore takes the additional signal of the past and uses it to transform the forecasting problem into a static one. This requires some weight placed on the revised signal, hence $K_r \in (0,1)$, and so some information from r_t is preserved in the forecast update. Therefore, the forecast update only contains learning on two shocks: w_t which is newly observed, and w_{t-1} which is now perfectly observed.[14]

After aggregating (1.7) I compute average forecast error to find $x_t - E_t[x_t] = (1-K)w_t$, and plug this plus the aggregate version of (1.8) into the CG regression (1.3):

$$(1-K)w_t = \beta(Kw_t + K_r w_{t-1}) + \varepsilon_t \tag{1.9}$$

Consider the true structural β when $\tau_v \to \infty$: $\bar{\beta}_{REV} = \frac{\rho - \rho K - K_r}{\rho K + K_r} = 0$. This means there is no relationship between forecast errors and forecast updates with a perfectly-revealing signal. Since the forecasting problem is static, forecast error only depends on w_t; but since w_t is not observed in the previous period, there is no information in the forecast update about it beyond what is in the current forecast. How one updates a static forecast across time periods thus has no predictability for their static forecast errors.

This result has implications when I turn to empirics, since many real-world data series are treated as if their revisions are perfect. For understanding the bias, it implies $\varepsilon_t = (1-K)w_t = \frac{K_r}{\rho}w_t$ as predicted from equation (1.4).[15] Estimating (1.9) by OLS gives the exact expression of the bias: $\hat{\beta} = \bar{\beta}_{REV} + \frac{(1-K)K}{K^2+K_r^2} = \frac{(1-K)K}{K^2+K_r^2}$. Due to the simple nature of the Kalman gains, I can

[14] I can rewrite (1.8) as $Kw_t + \rho(1-K)w_{t-1} + \tilde{\xi}_{it}$. The previous period's forecast put a weight of K on w_{t-1}, which is updated to $\rho(1-K)$ after learning r_t.

[15] If I take $K_r \to 0$, then $\bar{\beta}_{REV} = \frac{1-K}{K} = \bar{\beta}_{CG}$, and $\varepsilon_t = 0$ just as in CG. However, this is at odds with the assumption of $\tau_v \to \infty$, so is purely to demonstrate the coherence of these equations.

represent the bias in terms of fundamental parameters:

$$\text{Bias} = \frac{(1-K)K}{K^2 + K_r^2} = \frac{\tau_\psi \tau_w}{\tau_\psi^2 + \rho^2 \tau_w^2} \tag{1.10}$$

which is what (1.5) converges to as $\tau_v \to \infty$. The bias is strictly decreasing in the current signal precision; in the limit when $y_{it} = x_t$ the bias converges to 0, since the revised signal is no longer important. The bias is also decreasing in persistence, as demonstrated earlier in figure (1.1). Since the bias is the covariance between regressor and residual over the variance of the regressor, the presence of ρ is due to the variance of forecast updates. The forecast update is essentially an MA(1) in w_t, so increasing the persistence will always increase the variance of the process since past shocks will have a larger influence on the present.[16]

In terms of forecasts, increased persistence means the revised signal is more important, thus K_r will receive more weight. However, this does not affect the covariance between the residual and forecast update, since the revised signal does not help predict forecast errors. Instead, larger ρ increases the variance of each period's forecast, which translates into higher forecast update variance. In effect, the bias shrinks because forecasts follow the process and become more variable, which reduces the strength of connection between forecast errors and updates.

To summarize, when the revised signal is perfectly-revealing, forecasts no longer depend on prior expectations, which means forecast updates do not reflect updating of beliefs. Consequently there is no longer a relationship between forecast errors and forecast updates so now the regression residual equals the forecast error. There is still a covariance between forecast errors and forecast updates, but this is solely bias due to correlation between forecast errors and current period forecasts. The bias is decreasing in ρ because while the covariance does not depend on it, the variance does, which shrinks the whole term.

[16]Let $P_t = Kw_t + K_r w_{t-1}$. Define $\tilde{w}_t = \frac{w_t}{K}$, so $P_t = \tilde{w}_t + K_r \tilde{w}_{t-1}$. Increasing ρ increases K_r and not K, so this amounts to increasing the persistence of an MA(1).

Intuition: general revised signal

With a general revised signal $\tau_v \in (0,\infty)$ the intuition is similar, though now the true structural $\bar{\beta}_{REV}$ is non-zero. To see the additional persistence in the forecast update I rewrite the recursive forecasting equations as an ARMA representation. Repeatedly substituting equation (1.1) into itself gives:

$$\mathbb{E}_{it}[x_{t+h}] = \left(\frac{\rho^h(K+K_rL)}{(1-GL)}\right)x_t + \frac{\rho^h(K\psi_{it}+K_r\nu_t)}{1-GL} \tag{1.11}$$

where L is the lag operator. The extra dependence on the lagged state is immediate: the forecast has an additional moving average term due to the revised signal r_t. While the forecast made with one current signal is an AR(2) process in w_t, the forecast made with one current and one revised signal is an ARMA(2,1). Averaging (1.11) across agents and computing the forecast update in terms of the fundamental shock shows a similar equation to (1.8):

$$\mathbb{E}_t[x_{t+h}] - \mathbb{E}_{t-1}[x_{t+h}] = \rho^h \frac{Kw_t + K_r w_{t-1} + (1-\rho L)\nu_t}{1-GL} \tag{1.12}$$

The forecast update is now an ARMA(1,1) process in w_t, and includes updating the weight placed on the infinite sequence of past shocks. If $K_r = 0$ this becomes an AR(1) in w_t, exactly following the form of the underlying process: the forecast update reduces the weight on past shocks in a smooth manner, by the geometric sequence $\frac{1}{1-GL}$. With non-zero K_r there is an additional moving-average term that confounds this smooth discounting of the past. New information about the past adds a 'bump' to the flow of information when compared with the state process.

As before, I construct equation (1.3) when $h = 0$ by plugging in my expressions for the average forecast error and average forecast update processes:

$$\frac{(1-K)w_t - K_r\nu_t}{1-GL} = \beta\left(\frac{Kw_t + K_r w_{t-1} + (1-\rho L)K_r\nu_t}{1-GL}\right) + \varepsilon_t \tag{1.13}$$

Equation (1.13) shows how the same logic from the perfectly-revealing case still applies.[17] The forecast error only depends on an AR(1) in w_t, as dictated by the fundamental process. The forecast update is an ARMA(1,1) due to the revised signal. The extra weight on the past introduced by r_t creates a regression residual that is correlated with the regressor. Using $\bar{\beta}_{REV}$ and computing the residual finds the same terms identified in equation (1.4). If I assume ε_t is orthogonal and use OLS to find $\hat{\beta}$ I get:

$$\hat{\beta} = \frac{\tau_w^{-1}(1-K)(K+K_rG) - K_r^2\tau_v^{-1}(1-\rho G)}{\tau_w^{-1}(K^2+2KK_rG+K_r^2)+\tau_v^{-1}K_r^2(1-\rho^2+2\rho(\rho K+K_r))}$$

The difference between $\hat{\beta}$ and $\bar{\beta}_{REV}$ gives the same bias found in proposition 1. Note if $K_r \to 0$ (equivalently if $\tau_v \to 0$), $\hat{\beta} \to \bar{\beta}_{CG}$, so it is the revised signal which causes the bias. This expression is more complicated than in the perfectly-revealing case, and as such exhibits more complex comparative statics. Figure 1.2 plots bias equation (1.5) with a lower current signal precision ratio than in the previous figure. Overall the magnitude of the bias is larger than in figure 1.1 since there is less weight placed on the current signal. The bias still decreases in persistence, but for high enough ρ it becomes negative. As ρ increases, the influence of the fundamental falls due to increased variance of forecasts (similar to the perfectly-revealing case), but this is supplemented by the noise's negative contribution.

On surface level, the bias is due to representing the information learned from both current and revised signals by a single variable, the forecast update. Since the current signal contains w_t while the revised signal does not, a term in w_t shows up in the residual. Looking deeper into what constitutes the forecast error and forecast update, I showed that the revised signal introduces additional dependence on past shocks in the form of extra weight assigned to w_{t-1} and earlier. This extra weight does not affect the forecast error, so it shows up in the residual and thus biases estimates of $\bar{\beta}_{REV}$.

[17] Equation (1.13) is also the solution to the difference equation representation of forecast errors and forecast updates from section 1.2.1.

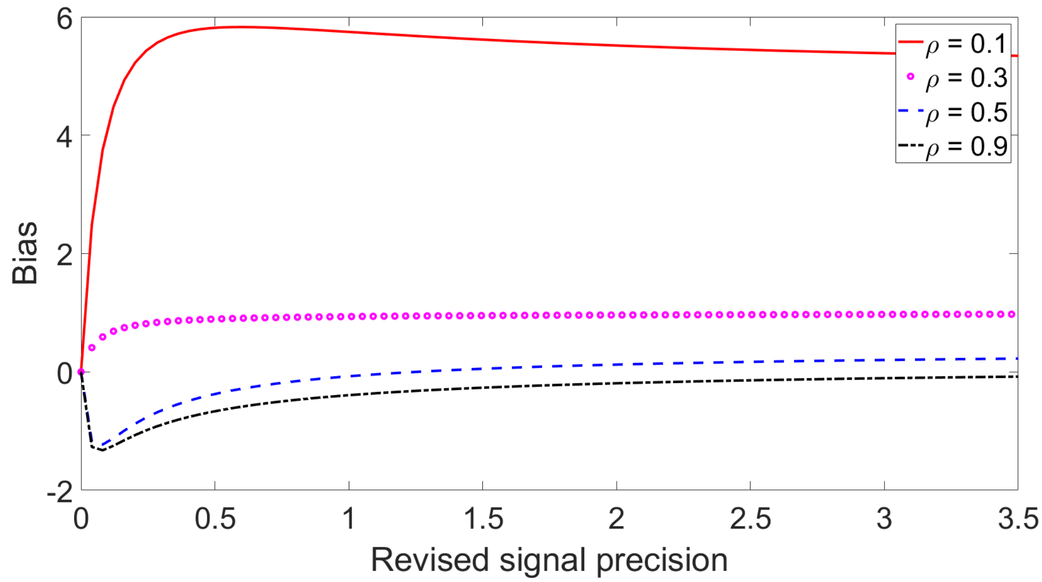

Figure 1.2. Bias when current signal precision is low. This graph plots the bias from equation (1.5) as the degree of revised signal precision changes for different levels of persistence ρ. The degree of current signal precision $\frac{\tau_\psi}{\tau_w} = 0.1$. When persistence is low (less than 0.5), the bias is positive and initially increasing until it begins to decrease at a low level of revised signal precision. When persistence is higher the bias is always negative and asymptotes to 0.

Next I discuss how an incorrect assumption of the structural β further biases the recovered degree of information frictions.

1.2.4 Implications for estimates of K

The implications for estimates of information frictions of the econometric bias crucially depend on the econometrician understanding the true structural β they are hoping to recover from (1.3) as seen in the mapping between $\hat{\beta}$ and K from equation (1.6). However, if the econometrician has a mistaken belief of the true structural β, then mapping $\hat{\beta}$ into K will introduce additional bias relative to that discussed previously. Specifically, the difference between the estimated $\hat{\beta}$ and the mistakenly-believed structural $\bar{\beta}_{CG}$ is the error in estimating K due to believing an incorrect structural relationship.

I first show this error with a perfectly-revealing revised signal:

$$\hat{\beta} - \bar{\beta}_{CG} = \frac{KK_r}{\rho(K^2 + K_r^2)} - \frac{1-K}{K} = -\frac{\rho^2 \tau_w^3}{\tau_\psi(\tau_\psi^2 + \rho^2 \tau_w^2)}$$

This shows that the recovered $\hat{\beta}$ will always be smaller than one expects it to be under the mistaken informational assumption of irrelevant revised data. When one believes the true structural β is $\bar{\beta}_{CG}$, then $\hat{K} = \frac{1}{1+\hat{\beta}}$. If $\bar{\beta}_{REV}$ is in fact the true structural β, then I can find the over-estimation of K by plugging $\bar{\beta}_{REV}$ into $\hat{\beta}$ and comparing $\hat{K}$ with the true K:

$$\hat{K} - K = \frac{\tau_\psi^2 + \rho^2 \tau_w^2}{\tau_\psi^2 + \rho^2 \tau_w^2 + \tau_\psi \tau_w} - \frac{\tau_\psi}{\tau_\psi + \tau_w} = K_r \left(\frac{\rho \tau_w^2}{(\tau_\psi^2 + \rho^2 \tau_w^2 + \tau_\psi \tau_w)} \right) \tag{1.14}$$

Expression (1.14) shows that mapping $\hat{\beta}$ into K as if there were no revised data will produce estimates of K that are too high, which means lower information frictions. This is the opposite effect on $\hat{\beta}$ and $\hat{K}$ compared with the econometric bias in the perfectly revealing case. Increased persistence causes estimates of K to be even larger than the truth.

To see the source of this error, return to equation (1.9). I showed that this equation cannot be estimated by OLS, since ε_t is not orthogonal to the forecast update. If one had the mistaken assumption that ε_t was orthogonal, that would imply $\beta = \bar{\beta}_{CG} = \frac{1-K}{K}$. Plugging this into (1.9) and solving for the econometric residual *under the belief* that $\hat{\beta} \to \bar{\beta}_{CG}$ gives:

$$-\frac{K_r}{K}(1-K)w_{t-1} = \varepsilon_t \tag{1.15}$$

If the revised signal were irrelevant, $K_r = 0$, and $\varepsilon_t = 0$ since $h = 0$. However, a useful revised signal causes the regression residual to be correlated with the regressor by $K_r w_{t-1}$, so OLS estimates will be biased. In this case, the correlation is through the lagged fundamental w_{t-1}, whereas the correlation for the econometric bias was due to w_t. The appearance of w_{t-1} is due to the assumption of $\beta \to \bar{\beta}_{CG}$, which ignores the extra persistence in forecasts from r_t.

1.2.5 Correct specification

The bias described above is essentially an omitted variable bias: the revised signal introduces an additional dependency in the forecast update which is not captured by the CG regression specification. To correct this I propose introducing an additional regressor to equation (1.3) that will capture the influence of the lagged observation. Since the revised signal is public, the econometrician has access to it, thus I can include the signal innovation term $(r_t - \mathbb{E}_{t-1}[x_{t-1}])$ in the specification to control for forecast updates due to the revised signal.

One can derive the following relationship from the optimal forecasting equation (1.1):

$$\begin{aligned} x_{t+h} - \mathbb{E}_t[x_{t+h}] = & \left(\frac{1-K}{K}\right)(\mathbb{E}_t[x_{t+h}] - \mathbb{E}_{t-1}[x_{t+h}]) \\ & - \rho^h \left(\frac{K_r}{K}\right)(r_t - \mathbb{E}_{t-1}[x_{t-1}]) + \gamma_{t+h} \end{aligned} \tag{1.16}$$

which suggests the regression specification:

$$x_{t+h} - \mathbb{E}_t[x_{t+h}] = \beta(\mathbb{E}_t[x_{t+h}] - \mathbb{E}_{t-1}[x_{t+h}]) + \beta_r(r_t - \mathbb{E}_{t-1}[x_{t-1}]) + \varepsilon_t \tag{1.17}$$

The first term of equation (1.16) as well as the true structural β is the same as in (1.2). The second term represents the information learned from signal r_t. The residual $\varepsilon_t = \gamma_{t+h}$ no longer contains the fundamental shock w_t and contains no noise terms since they are preserved in the signal r_t, so this equation can be consistently estimated by OLS. Estimates of β and β_r directly map into K and K_r and do not require a measure of ρ.

The structural relationship (1.16) implies several testable restrictions. First, it predicts the coefficient β_r should be negative: since observing r_t doesn't affect forecast errors but does affect forecast updates, $\beta_r < 0$ in order to remove the impact of r_t from the forecast update.[18]As the

[18]To see this, return to the perfectly-revealing case: $(1-K)w_t = \frac{1-K}{K}(Kw_t + K_r w_{t-1}) - \frac{K_r}{K}(r_t - \mathbb{E}_{t-1}[r_t]) + \varepsilon_t$. Plug in $K_r = \rho(1-K)$ and $\mathbb{E}_{t-1}[r_t] = \rho x_{t-2} + Kw_{t-1} + \xi_{t-1}$, simplify to find $(1-K)w_t = (1-K)w_t + \rho\frac{(1-K)^2}{K}w_{t-1} - \rho\frac{(1-K)^2}{K}w_{t-1} + \tilde{\xi}_t + \varepsilon_t$.

horizon increases, β_r will decrease, but similar to CG β should remain constant across horizons. Note that any prediction across variables is conditional on constant precisions τ_ν and τ_ψ, plus constant variance of the fundamental τ_w^{-1}, which is unlikely to be true in the data. The predictions within a variable are more likely to hold empirically.

1.3 Data

To assess the validity of the model I require a panel of forecasts to construct average forecasts, as well as real-time data releases in order to construct the revised signal innovation. Here I describe the data sources, paying special attention to the schedule of data releases and sources of revisions.

1.3.1 Expectations

I use the publicly available Survey of Professional Forecasters (SPF) for expectations data. The survey includes a panel of roughly 30-40 forecasters making quarterly predictions on key economic indicators for horizons up to four quarters ahead. The SPF also includes a forecast of the previous quarter, which I will show is necessary for estimating certain data series. Most forecasters use a mathematical model, adjusted with subjective judgments, to construct their forecasts, while 60% update forecasts at a monthly or higher frequency (Stark, 2013). This suggests the SPF forecasters are highly attentive to new data releases so their forecasts will likely respond to revised signals, which makes the data ideal to test the model.

The survey is sent to respondents at the end of the first month of each quarter, timed to arrive after the advance report of the national income and product accounts (NIPA) by the Bureau of Economic Analysis. The deadline for returning the survey varies around the middle of the second month of the quarter, although since 2005 the deadline is by the second week.[19] The information set for respondents thus includes the advance NIPA report, anything released before the middle of the second month, and everything published in the past. For variables published

[19]The Philadelphia Fed lists the exact deadline dates on their website.

quarterly respondents should only receive one current signal (e.g. the advance report for NIPA variables) per survey. For variables published monthly respondents could receive an additional signal during the second month before the SPF deadline.

1.3.2 Real-time Macroeconomic Data

Most macroeconomic data series are not perfect representations of the true underlying value, but instead are estimates based on statistical sampling. The first estimate of a variable rarely has access to all the underlying raw data, which are collected later and incorporated into subsequent estimates. I describe the revisions process here for two sets of variables: those from the National Income Product Accounts (NIPA) and those from other sources.

NIPA variables

The BEA releases its advance report of the NIPA at the end of the first month of each quarter. This report has the first quarterly estimates of nominal and real GDP, the GDP price index, real consumption, real investment, etc. The second and third estimates are released at the end of the second and third months of each quarter.[20] This never occurs before the SPF deadline, so only the advance estimate is in the forecaster's information set. Additionally, there are yearly revisions to all quarterly estimates of the preceding three calendar years, plus a comprehensive revision every five years based on economic censuses. Since the annual and comprehensive revisions often include definitional changes I always compare values within a data vintage when, for example, computing growth rates.

The difference between the advance and second estimates of the NIPA is due to the different source data available. The NIPA handbook describes five categories of source data which can be ranked according to their comprehensiveness and accuracy. Holdren (2014) shows that the most comprehensive data only comprise 25% of the advance estimate, while trend calculations (based solely on past observations) contribute 28%. One month later, the second

[20] Appendix A.1.3 demonstrates how including additional revised signals in the model does not qualitatively change the results, though now there are additional biases due to each revision.

estimate uses 37% comprehensive data and only 20% trend. The third estimate uses the least trend data, with other more comprehensive categories picking up the difference.

To be more specific, the advance estimate has no source data available for the third month of the target quarter for many key series, like construction spending, changes in inventories for nondurable-goods, and U.S. exports and imports of goods (Holdren, 2014). The BEA makes trend-based assumptions for these series in the third month based on past periods. The second estimate replaces a statistical prediction of, for example, private construction spending, with source data from the census bureau's monthly Value of Construction Put in Place Survey. This report is released one month after the target month, so it would be unavailable for the third month of the advance NIPA estimate target quarter. Another example is the Quarterly Services Report (QSR) used to calculate private investment and PCE inflation. The QSR releases 2.5 months after the target month, so can only be incorporated in the second (and sometimes third) NIPA estimate. In sum, the revisions to the NIPA represent increased knowledge of the underlying variable through more comprehensive and representative data availability.

Importantly, the NIPA reports contain estimates only of the prior quarter: the advance report released in Q1 2022 has the first estimates for Q4 2021. Appendix A.1.4 shows an alternative way of modeling the NIPA data which incorporates the delay in data publications. Empirical results are similar across both methods, so I will use the initial model presented in section 1.2.

Labor statistics

The Bureau of Labor Statistics publishes the Employment Report every month on the first (or second) Friday.[21] The report contains the results from two surveys: the Current Population Survey (a survey of households) and the Current Employment Statistics (a survey of establishments). These surveys provide the first monthly estimate of the unemployment rate (from the CPS) and nonfarm payroll employment (from the CES) for the previous month.

[21]Technically: 'The Employment Situation is scheduled for the third Friday following the week including the 12th of the prior month, with an exception for January.'

The CES is a voluntary survey collected from firms about their employee levels, hours and earnings. As a voluntary survey, the CES fails to receive 100% of sampled firms responses by the publication date of the Employment Report. In fact, the average response rate by the first publication of the target month over the last 40 years is around 61%.[22] The BLS continues to collect responses after the first Employment Report, which they call the first preliminary release, incorporating the additional survey data into revised estimates in the next two Employment Reports. Thus the revised monthly estimates are due to additional survey receipts, which should improve the accuracy of the estimates.

The monthly unemployment rate from the CPS is not revised on a month-to-month basis, as the survey completes its data collection prior to publication. The BLS does annually adjust the population weights of the survey according to census data, but this is done in real-time and does not affect historic data. Therefore, a second estimate of the monthly unemployment rate would be identical to the first estimate.

There is an additional dimension of data revision to both nonfarm payroll employment and the unemployment rate. In order to match the SPF forecasts, I take quarterly averages of the monthly employment rates. The consistent timing of the SPF and the Employment Report means a second estimate of the quarterly employment level or unemployment rate would be different from the first estimate due to additional months of data.

As an example, consider the SPF for 2019 quarter 3. The survey is sent to respondents at the end of the first month, July. The June Employment Report, containing the unemployment rate in June, was published on the 5 July. The information set of forecasters contains the unemployment rates for each month in quarter 2: April, May, June. The deadline for returning the survey is 6 August. On 2 August, the next Employment Report is released and reveals the unemployment rate in July, the first month of quarter 3. Forecasters take this signal as the average unemployment rate in quarter 3 despite it only representing July. Fast forward to the next SPF for quarter 4. By this survey deadline, forecasters will have received the September, October,

[22]This has trended up over time as the BLS improves its electronic data collection systems.

and November Employment reports with August, September and October unemployment rates. Now consider their information sets across both surveys. For the quarter 3 survey, all they knew of the quarterly unemployment rate in quarter 3 is the July monthly figure. For the quarter 4 survey, they've learned the August and September figures as well. I consider this addition of new data as the quarterly revision, despite none of the monthly unemployment rates changing. The employment levels will change, which means this variable has two dimensions of revision. Furthermore, this means that BLS variables are available for the current time period, even if it's only one month out of the quarter.

Financial variables

I study three financial variables: the secondary market interest rate on Treasury bills (3 month) and bonds (10 years) and Moody's AAA corporate bond yield index. These variables never receive revisions like the NIPA data do, and they can be publicly looked up at any given moment. However, since the SPF asks respondents for their forecast of the quarterly average interest rate, their initial forecast will be based on incomplete information of the quarter. When an individual makes a forecast in Q1 of the quarterly average Treasury bill rate for Q1, for example, they only have access to the daily rates up through the middle of February. In effect, their nowcast is a prediction of the second half of Q1. Next quarter when they make forecasts they will now know the true average rate in Q1. Therefore, I use the average rate up to the deadline date for returning the survey each quarter as the initial signal y_{it}, and the complete average rate as known on the first day of the following quarter as the revision r_t.

Housing starts

The U.S. Census Bureau publishes a monthly report on New Residential Construction (NRC), where the number of housing starts is collected in the Survey of Construction (SOC). The data are collected by Census field representatives visiting permit offices and building sites to catalog which permits have begun construction. Difficulties in reaching every site or permitowner in the sample means the first estimate of housing starts is based on incomplete data. The Census Bureau releases revisions to their first estimate in the following two months as the sample becomes fully covered.

Published estimates of monthly housing starts correspond to the previous month. Unlike the Employment Report, the NRC is published during the third week of each month, and thus will be after the SPF deadline. Only the first month of each quarter is in forecasters' information set. Therefore the only source of revisions to housing starts is the receipt of additional source data.

Industrial production index

The Federal Reserve Board of Governors publishes the G.17 report monthly which contains their estimates of real industrial output and capacity utilization. There are two main sources of data: physical product output and production inputs. Physical output data are received from private trade agencies at a monthly or quarterly frequency, with some quarterly data receipts coming several months after the initial estimate. When product output data are not available, the Board of Governors converts estimates of production-worker hours from the BLS establishment survey into production outputs using factors derived from historical relationships. When some of either source data is not yet available the Board of Governors estimates the missing production using statistical modeling. In 2020, the first monthly estimate of the index was derived from 31% physical product data, 43% production-worker hours, and the remaining 26% of the index was estimated. By the third monthly estimate source data accounted for 93%. The sixth and final monthly estimate had 98% source data. The first estimate is released around the 15th of the following month, which means the forecaster's information set contains the first estimate of the first month of the current quarter.

1.3.3 Data caveats

The first concern is the mismatched frequency of data revisions and forecast updates. Since the SPF is conducted quarterly but most data series are revised monthly, I do not observe forecast updates in response to the initial revisions. For example, the advance estimate of Q1 is released during the first month of Q2 and is available for the Q2 forecaster survey. After forecasts are submitted, the second estimate (i.e. first revision) will be released during the second month of Q2. Before the next SPF occurs the third estimate will be released (i.e. second revision), as well as possibly annual or benchmark revisions. This means the forecast update I observe in the Q3 SPF is reacting to the third or fourth revision of Q1. The size of the revision from the advanced to the third estimate is larger than from the advanced to second estimate which may or may not improve the empirical identification.

The lagged release schedule of the NIPA variables presents its own problems. Many forecasters use the advance estimate as their backcast $\mathbb{E}_{it}[x_{t-1}]$. Using the model from section 1.2, this implies forecasters believe r_t is perfectly-revealing. Table (1.1) summarizes the share of individual backcasts made that equal the advance estimate. The high shares across NIPA variables indicate that I should view the empirical results as if $\tau_v \to \infty$. In that case, the true structural $\bar{\beta}_{REV} = 0$, so any statistical evidence to the contrary is due to bias.

Table 1.1. Share of Individual Backcasts Equaling Real-time Signal

NGDP	**RGDP**	**Price Index**	**Non-Resid. Inv.**	**Resid. Inv.**	**Cons.**
90%	82%	86%	56%	54%	61%

Notes: Each column computes the share of individual backcasts (i.e. forecasts of x_{t-1}) that are within 0.001 of the advance release of that data series. An error margin of 0.001 is used to include cases where forecasters appeared to round the data release differently. Data are from 1968Q4 to 2020Q3. The Price Index data series is the GDP deflator.

1.4 Empirically testing the model

I wish to test the predictions of the model with a revised signal and compare the estimates of information frictions to the standard specification of CG. As a first step, I perform a test of the relationship between data revisions and forecast updates. If forecasters pay attention to revised data, there should be a correlation between the size of data revisions and the forecaster's own update.[23] I use the notation $x_{t-s|t+j}$ to denote the value of x_{t-s} as published in $t+j$. For example, from the model $y_{it} = x_{t|t} + \psi_{it}$ and $r_t = x_{t-1|t} + \nu_t$. I define the size of the data revision observed in period t as the difference between the initial release and the first revision. For non-NIPA variables, this is $z_t^O = x_{t-1|t} - x_{t-1|t-1}$. Data revisions to the NIPA are $z_t^N = x_{t-2|t} - x_{t-2|t-1}$ due to the lagged release schedule. I run the following regression:

$$\mathbb{E}_t[x_{t+h}] - \mathbb{E}_{t-1}[x_{t+h}] = c + \alpha_1 z_t^d + \delta_t \tag{1.18}$$

The null hypothesis is $\alpha_1 = 0$, implying forecasters do not condition their forecast updates on revised information. Table (1.2) shows results when the forecast horizon is one year.[24] I cannot reject the null for most of the NIPA variables plus housing starts and industrial production, yet the coefficient on nearly all other non-NIPA variables is statistically different from 0 at the 10% level. The discrepancy is likely due to the delayed release schedule of the NIPA variables. Recall from table 1.1 that most forecasters use the current signal as their backcast. This implies they believe the current signal is the truth (or at least close-enough to it), so they have no reason to monitor revisions. On the other hand, the results for non-NIPA variables suggest forecasters do pay attention to revised data signals. The advance estimate for these data series are published in the current period but only contain one month of the quarter. This makes the second estimate (i.e. the data revision) important for accurate forecasting.

[23] This should hold at an individual level as well, but I focus on the aggregate forecast update to stay in line with the model.

[24] Appendix A.1.6 contains results for other horizons, as well as including the first data release as a control in the regression. Results are similar across horizons and with the control.

Table 1.2. Data Revisions Predicting Aggregate Forecast Updates, one-year ahead

Forecast Updates	Data Revision	pval
Nominal GDP	-0.08	0.75
Real GDP	0.06	0.75
GDP price index inflation	0.32	0.53
Real consumption	**0.34**	0.02
Real nonresidential investment	0.04	0.75
Real residential investment	0.04	0.82
Consumer price index	**0.16**	0.00
Housing starts	0.06	0.46
Industrial production	0.03	0.18
Unemployment rate	**1.34**	0.00
Employment growth	**1.33**	0.00
Treasury bill 3 month rate	**1.41**	0.00
Treasury bond 10 year rate	**1.76**	0.00
AAA Corporate bond rate	**1.06**	0.00

Notes: Each row shows the estimated coefficient $\hat{\alpha}_1$ from regression model (1.18) for one-year ahead forecasts. The variables in the top panel are from the NIPA. Data used for estimation are from 1968Q4 to 2019Q4. Standard errors are computed using the Newey-West method with automatic bandwidth selection to correct for serial correlation. Estimates are bold if I reject the null hypothesis of ignoring revised data releases at the 10% level.

Importantly, there is positive correlation between forecast updates and data revisions for the non-NIPA variables. This implies that a revision which increases the value of x_{t-1} will lead to an increase in the forecast update. Put another way, an unexpectedly high revision causes forecasters to raise their current forecast of x_t. These preliminary results suggest I am more likely to find validation of the model in the non-NIPA variables than the NIPA variables.

1.4.1 Main results

To test the validity of the model, I estimate (1.17) using the third estimate (second revision) as x_t (i.e. $x_{t|t+3}$). Table 1.3 presents regression results for one-year ahead forecasts in both the CG model without revisions (equation 1.3) and the corrected specification.

Table 1.3. Average Forecast Errors and Average Forecast Updates, one-year ahead

	CG Model			**Revised Model**				
Variable	β	**pval**	R^2	β	**pval**	β_r	**pval**	R^2
Nominal GDP	**0.68**	0.00	0.05	**0.72**	0.00	**-0.53**	0.02	0.07
Real GDP	**0.63**	0.02	0.04	**0.66**	0.02	-0.37	0.11	0.05
GDP price index inflation	**1.35**	0.01	0.24	**1.28**	0.01	**0.63**	0.00	0.26
Real consumption	0.36	0.25	0.01	0.29	0.38	**0.33**	0.06	0.02
Real nonresidential investment	**0.97**	0.00	0.07	**0.90**	0.02	0.30	0.45	0.07
Real residential investment	**1.17**	0.00	0.10	**1.09**	0.00	**0.49**	0.03	0.11
CPI inflation	0.24	0.30	0.00	**0.74**	0.00	**-0.25**	0.00	0.05
Housing starts	0.42	0.11	0.01	**1.35**	0.00	**1.36**	0.00	0.17
Industrial production	**0.95**	0.00	0.08	**0.94**	0.00	**-0.67**	0.02	0.10
Unemployment rate	**0.83**	0.00	0.12	**1.17**	0.00	**-1.08**	0.01	0.15
Employment growth	**1.60**	0.00	0.34	**1.49**	0.00	**1.39**	0.00	0.36
Treasury bill 3 month rate	**0.67**	0.00	0.08	**0.26**	0.08	**1.97**	0.00	0.21
Treasury bond 10 year rate	-0.06	0.42	-0.01	-0.05	0.72	**-0.55**	0.06	0.00
AAA Corporate bond rate	0.08	0.61	-0.01	**-0.30**	0.02	**1.30**	0.01	0.11

Notes: The CG model refers to estimates of (1.3), while the Revised Model refers to estimates of (1.17). All forecasts are of x_t one-year ahead. The variables in the top panel are from the NIPA. Data used for estimation are from 1968Q4 to 2019Q4. Standard errors are computed using the Newey-West method with automatic bandwidth selection to correct for serial correlation. Estimates that are in bold are statistically significant at the 10% level.

Recall that β is the coefficient on the average forecast update and β_r is the coefficient on the correction term, the revised signal innovation. Testing the model's key assumption that forecasters use revised data has a null hypothesis of $\beta_r = 0$. The data reject at the 10% level for all variables excluding real GDP and nonresidential investment. This is strong support of the model, as CG point out that no other variable should have predictive power once conditioning on forecast updates. Additionally, the revised model finds a statistically significant β for nearly all variables, while the CG model does not for CPI inflation, housing starts, and the corporate bond rate. This means that estimates from CG would predict no information frictions on those variables. Finally the model fit as measured by R^2 is much greater for the revised model when considering non-NIPA variables: R^2 for 3-month treasury bills increases from 8% to 21%, for example. The fit is mostly the same for NIPA variables, which may reflect the results from table 1.2.

Next I compute the influence of the revised data signal on the recovered estimate of information frictions. Table 1.4 performs a statistical comparison of $\hat{\beta}$ between the two models and computes the implied estimate of K. The last column takes the difference and is a measure of the bias from using the CG model when the true information set contains revised data. As hinted at in earlier results, none of the NIPA variables have a statistically distinguishable estimate of K across the two models. Despite the revision correction term being significant, the resulting estimate of K is unchanged. At the same time, five of the eight non-NIPA variables do have a significantly different estimate of K. Averaging the absolute difference across these five variables finds an average bias of 0.26, or 0.17 when averaged across all eight. The direction of the bias varies: for non-financial variables it is positive, while bond rates have a negative bias.

Table 1.4. Comparing estimates of K between CG and Revised model, one-year ahead

Variable	Test stat	pval	CG K	REV K	Difference
Nominal GDP	0.89	0.34	0.59	0.58	0.01
Real GDP	0.54	0.46	0.61	0.60	0.01
GDP price index inflation	1.32	0.25	0.43	0.44	-0.01
Real consumption	2.40	0.12	0.73	0.78	-0.04
Real nonresidential investment	0.61	0.44	0.51	0.53	-0.02
Real residential investment	1.41	0.23	0.46	0.48	-0.02
CPI inflation	13.03	0.00	0.80	0.57	**0.23**
Housing starts	21.87	0.00	0.70	0.43	**0.28**
Industrial production	0.06	0.81	0.51	0.52	0.00
Unemployment rate	5.17	0.02	0.55	0.46	**0.08**
Employment growth	1.55	0.21	0.38	0.40	-0.02
Treasury bill 3 month rate	8.25	0.00	0.60	0.79	**-0.20**
Treasury bond 10 year rate	0.36	0.55	1.07	1.05	0.02
AAA Corporate bond rate	4.62	0.03	0.93	1.42	**-0.49**

Notes: The test statistic is a chi-squared adjusted for heteroskedasticity. CG K and REV K refer to the K recovered from both model estimates according to the expression $K = \frac{1}{1+\hat{\beta}}$. Estimates that are in bold are statistically different at the 10% level.

In appendix A.1.7 I show versions of tables 1.3 and 1.4 across all horizons. The same variables exhibit large bias across horizons, though the size of this bias varies greatly. This is contrary to the model's predictions, but it affects the CG specification similarly.

1.4.2 Testing the relationship at the individual level

Bordalo et al. (2020) estimate the CG regression (1.3) at an individual level and find evidence of over-reaction to data, or $\hat{K} > 1$. According to rational expectations there should be no relationship between individual forecast errors and individual forecast updates, since if there were the forecaster could exploit it and improve their forecast. Finding $K > 1$ casts doubt on the rationality of forecasters and suggests individuals suffer from a behavioral bias. Ryngaert (2017) found some evidence of overestimation on the individual level in a model of mis-perception of the fundamental persistence.

In this section I will re-run my corrected specification at the individual forecaster level to assess if controlling for revised data can address these findings. Since revised data can introduce a positive or negative bias into estimates of K, depending on parameters, it is possible that the results in Bordalo et al. (2020) and Kohlhas and Walther (2021) can be explained in terms of the identified bias. Following these papers I pool across forecasters and recreate tables 1.3 and 1.4 at the individual level.

Table 1.5. Average Forecast Errors and Average Forecast Updates
Individual level, one-year ahead

	CG Model			**Revised Model**				
Variable	β	**pval**	R^2	β	**pval**	β_r	**pval**	R^2
Nominal GDP	**-0.23**	0.01	0.02	**-0.28**	0.00	**0.34**	0.01	0.04
Real GDP	**-0.17**	0.08	0.01	**-0.14**	0.23	**0.33**	0.17	0.01
GDP price index inflation	-0.09	0.60	0.00	-0.17	0.14	1.17	0.00	0.09
Real consumption	**-0.22**	0.05	0.02	**-0.23**	0.04	**0.34**	0.04	0.04
Real nonresidential investment	0.04	0.80	0.00	-0.05	0.78	0.74	0.03	0.05
Real residential investment	-0.03	0.79	0.00	-0.04	0.73	0.79	0.00	0.06
CPI inflation	**-0.27**	0.02	0.02	**-0.30**	0.00	**0.00**	0.95	0.02
Housing starts	**-0.29**	0.01	0.02	**-0.04**	0.69	**0.84**	0.00	0.12
Industrial production	-0.04	0.69	0.00	-0.04	0.61	0.27	0.28	0.01
Unemployment rate	**0.33**	0.04	0.03	**0.28**	0.07	**0.22**	0.37	0.03
Employment growth	0.36	0.25	0.03	0.12	0.63	1.70	0.00	0.12
Treasury bill 3 month rate	**0.27**	0.01	0.02	**0.03**	0.69	**1.50**	0.00	0.16
Treasury bond 10 year rate	**-0.19**	0.02	0.01	**-0.33**	0.00	**0.54**	0.00	0.04
AAA Corporate bond rate	**-0.21**	0.00	0.02	**-0.40**	0.00	**0.90**	0.00	0.12

Notes: The CG model refers to estimates of (1.3), while the Revised Model refers to estimates of (1.17). The table pools all individual forecasts together. All forecasts are of x_t one-year ahead. The variables in the top panel are from the NIPA. Data used for estimation are from 1968Q4 to 2019Q4. Standard errors are clustered by forecaster and quarter, and corrected for serial correlation and heteroskedasticity with a bandwidth of four quarters. Estimates that are in bold are statistically significant at the 10% level.

Table 1.6. Comparing estimates of K between CG and Revised model
Individual level, one-year ahead

Variable	Test stat	pval	CG K	REV K	Difference
Nominal GDP	18.29	0.00	1.31	1.40	**-0.09**
Real GDP	2.63	0.10	1.20	1.16	**0.04**
GDP price index inflation	4.13	0.04	1.10	1.20	**-0.10**
Real consumption	0.16	0.69	1.29	1.30	-0.01
Real nonresidential investment	5.88	0.02	0.96	1.05	**-0.09**
Real residential investment	0.11	0.73	1.03	1.04	-0.01
CPI inflation	0.28	0.60	1.36	1.43	-0.06
Housing starts	34.15	0.00	1.41	1.04	**0.36**
Industrial production	0.37	0.54	1.04	1.04	-0.01
Unemployment rate	0.96	0.33	0.75	0.78	-0.03
Employment growth	11.45	0.00	0.73	0.89	**-0.16**
Treasury bill 3 month rate	32.11	0.00	0.79	0.97	**-0.18**
Treasury bond 10 year rate	9.83	0.00	1.23	1.48	**-0.25**
AAA Corporate bond rate	33.71	0.00	1.27	1.68	**-0.40**

Notes: The test statistic is a chi-squared adjusted for heteroskedasticity. CG K and REV K refer to the K recovered from both model estimates according to the expression $K = \frac{1}{1+\hat{\beta}}$. Estimates that are in bold are statistically different at the 10% level.

At the individual level there is still strong evidence in favor of forecasters using revised data, since I can reject $\beta_r = 0$ for the majority of variables. Most estimates of β are now negative across both models, which matches the results from Bordalo et al. (2020). Table 1.6 shows how the revision correction alters estimates of K at the individual level. Although nine of the thirteen variables have different estimates of information frictions in the revised model, there is rarely a difference in implication. Variables that have over-reaction continue to do so, with housing starts being a possible exception. Variables that have under-reaction largely remain so, except for the three-month treasury bill, which the revised model predicts has no information friction. In total, the revision correct does not reverse findings of individual over-reaction; if anything, the correction strengthens them.

1.5 Applications: Model Validation and Forward Guidance

This section demonstrates the economic significance of my empirical findings that information frictions are under-estimated by 26 percentage points on average. For the first application, I take a theoretical New Keynesian model with incomplete information and show how the implied impulse response function better matches the data when using my estimate of K compared with the CG estimate. In the second application, I use a model of forward guidance to show that with my estimate of K the policy has severely diminished efficacy. The goal is to highlight the importance of correctly measuring information frictions by using leading models from the literature. It is important to note that these models do not include a revised signal in their information structure, so the following analysis is only illustrative. To fully study how a revised signal interacts within an incomplete information model requires a new theory which embeds my information structure.

1.5.1 Impulse responses in the New Keynesian model

Angeletos and Huo (2021) develop a general theory which links incomplete information and two behavioral distortions. They show that their theory can nest two components of the New

Keynesian (NK) model: the investment-saving curve and the Phillips curve (NKPC). For the latter, they assume an exogenous AR(1) process for marginal cost which firms observe noisily. They solve the model using the Wiener-Hopf technique, first used in an economics context by Rondina (2008), and then study how inflation responds to marginal cost shocks. I recreate their impulse response functions in figure 1.3.

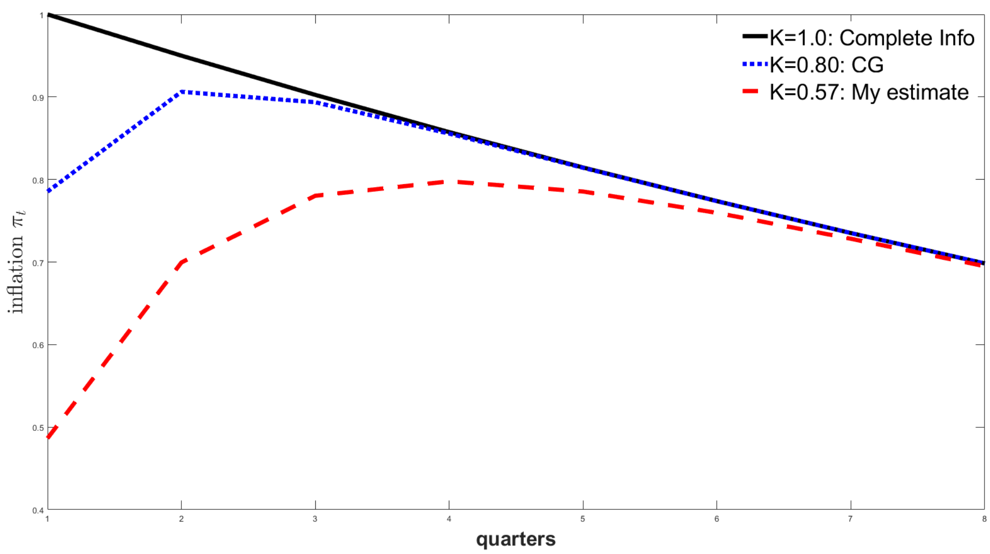

Figure 1.3. The IRFs plotted above come from the hybrid NKPC in Angeletos and Huo (2021). The black solid line is when information is complete. The dotted blue line uses an estimate of information frictions according to the CG specification, while the dashed red line uses the estimate from the revised model. Both estimates are from CPI inflation.

Following a marginal cost shock, inflation jumps to its highest level then smoothly returns to its previous value when information is complete. This is at odds with empirically estimated IRFs as in Gali, Gertler and David Lopez-Salido (2005), which demonstrate a hump-shaped profile. Relaxing the complete information assumption by introducing noise into observations of marginal cost brings the IRF closer to the data. The dashed blue line uses the estimate of K from the CG model applied to CPI inflation. The initial jump is not as high as the complete information benchmark, and inflation continues to rise for one additional period before following the smooth discounting of the complete information curve.

Contrast this with the dashed red line, which uses the estimate of K from the revised model. The initial response to the marginal cost shock is markedly modest, and it takes a full year before inflation reaches its peak response. This better matches the evidence from the VAR literature of weak immediate reactions followed by prolonged build-up. The difference between the blue and red lines matters for how successful incomplete information theories are at matching the data. With only the blue line, one may believe information frictions are insufficiently strong to induce the degree of persistent witnessed in the data. Instead, the red line suggests incomplete information can be the sole friction needed to achieve this persistence. Generating an accurate estimate of information frictions is vital for proper model validation.

1.5.2 Efficacy of forward guidance

Forward guidance is a monetary policy tool which attempts to sway expectations of future monetary action through statements of commitment. Angeletos and Lian (2018) study the power of this policy action through the lens of an incomplete information model. The main conclusion they reach is that without common knowledge, general equilibrium effects (through higher-order beliefs) greatly attenuate the impact of future monetary policy commitments, which contradicts the prediction of a sustained response when information is complete. They illustrate this fact by plotting, for increasing horizons, the ratio of consumption response to forward guidance with incomplete information to the response with complete information. A value of 1 for this ratio implies the response is invariant to information, while a value of less than 1 implies incomplete information reduces the efficacy of the announcement.

I recreate this attenuation curve for two values of K: 0.8 as estimated from the CG model, and 0.58 as estimated from the revised model. The results are in figure 1.4. The dashed blue line captures the attenuation result from their paper: consumption response falls as the periods since the announcement increase. After four years, the response is less than 25% that of the complete information benchmark. The dotted red line shows the degree of attenuation with the revised model estimate of $K = 0.57$. This curve shows a much steeper initial decline in efficacy, such that the response at four years is under 5% the benchmark.

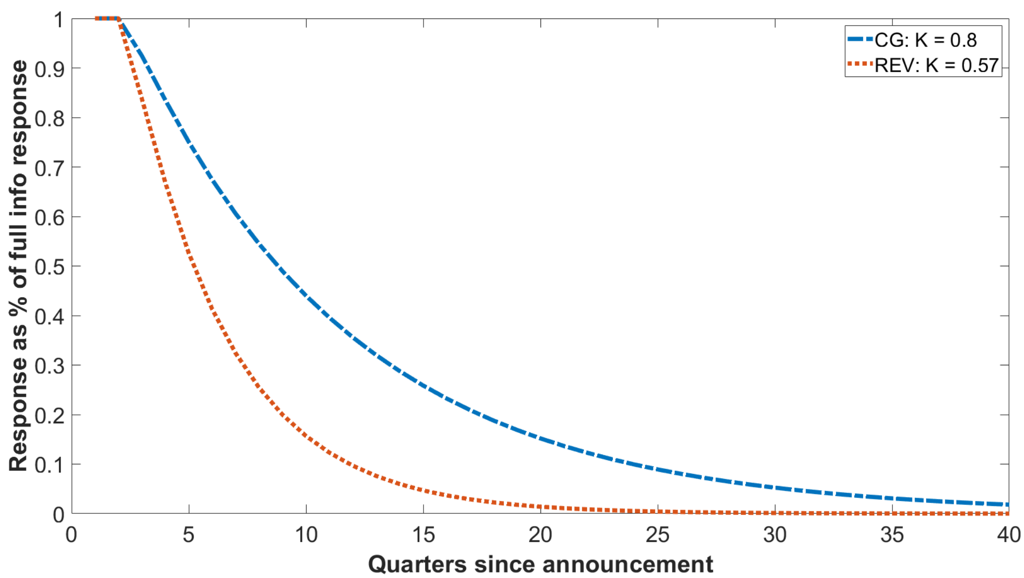

Figure 1.4. Consumption response to forward guidance announcement. The curves plotted above come from the incomplete information model of Angeletos and Lian (2018). The dashed blue line is when the degree of information frictions is 0.8, as estimated from the CG model. The dotted red line uses an estimate of 0.58 from the revised model. The y-axis is the ratio of consumption spending in response to a forward guidance announcement relative to the same response when information is complete.

Both curves showcase the ability of information frictions to solve the forward guidance puzzle, though the red line offers a more pessimistic view of the policy's usefulness. Determining which line is correct is important when considering what mixture of unconventional monetary tools to use when nominal interest rates are stuck at the zero lower bound.

1.6 Conclusion

Information frictions can be a powerful tool for studying policy and improving model fit, but this crucially depends on what degree of friction we impart to the model. I have shown that our ability to estimate information frictions from survey data is a function of the information structure we assume survey participants are using. When revised data signals are present, rational forecasts have additional dependence on past outcomes which biases the standard econometric specification between forecast errors and updates. Without correction, regression estimates no longer recover the degree of information frictions. I propose a model-implied correction term to capture the extra persistence in forecasts which fully removes the bias. My new estimates suggest information frictions are 20 percentage points lower on average across three key indicator variables (CPI inflation, housing starts, unemployment), with an absolute average difference of 26 percentage points when including financial variables. Applying my regression model to test individual rationality yields noisy results which do not overturn existing ideas in the literature. Despite this, I find strong evidence across specifications that revised data are an important signal for professional forecasters; future empirical work should consider revised data when deriving new test specifications.

Chapter 1 is currently being prepared for submission for publication; Tyler L. Paul. The dissertation author was the principal author on this paper.

Chapter 2

Now or Later: Optimal Data Provision with Data Revisions and Strategic Incentives

When statistical offices (S.O.) are budget constrained they face a trade-off: provide imprecise yet timely data about the present, or wait to provide more precise yet delayed data about the past. I study this decision in the context of a beauty-contest economy where end-users of data want to produce forecasts that are both accurate and similar to one another. The S.O.'s welfare-maximizing choice of current versus revised signal precisions depends on the accuracy of private information and strength of strategic incentives. The S.O. prioritizes the revised signal when individuals have poor private information, yet as it becomes more precise the S.O. increases the budget devoted to the current signal until eventually it does not provide a revised signal. When private information and strategic incentives are sufficiently high, this reverses: the S.O. increases resources towards the revised signal in order to reduce the crowding-out of private information caused by the current signal. A larger budget constraint makes the S.O. reduce the precision of the revised signal in favor of the current signal, as the S.O. effectively treats the former as an inferior good.

2.1 Introduction

The role of a national statistical office (S.O.) is to publish estimates of nation-wide aggregate variables in a timely manner. In an ideal world, the S.O. could perfectly observe all economic activity over a given time period so that their estimates comprehensively and accurately reflect true activity. In reality the S.O. has a limited budget for data collection and thus relies on surveys of a representative group of economic actors. Conducting and analyzing these surveys takes time and labor resources, so when the data publishing deadline approaches the S.O. typically has an incomplete set of sample data to use when estimating the national aggregate. As a result, the initial value of a variable is usually revised in later periods due to the receipt of more survey responses.

A limited budget combined with the nature of data collection implies the S.O. must trade-off the timeliness of estimates with their accuracy. To understand the consequences of this trade-off the S.O. needs to consider the behavior of their end-users, and assess the marginal value of more accurate data releases in the next period versus the value of timely data today. That is what this paper addresses: how should the S.O. use its limited resources to provide data releases that maximize social welfare? The answer to this question can provide a set of guidelines for a S.O. to follow according to the unique features of their economy and budgetary environment.

I begin by casting the question into a two-staged optimization problem. The first stage involves the end-users of data solving a forecasting problem in which they use noisy signals to simultaneously minimize their forecast error as well as the distance between their guess and the average guess. The latter objective reflects strategic complementarities, an additional incentive for agents to coordinate their forecasts along the lines of Keynes's beauty contest. This model belongs to a class of economies called global games which offer a flexible and general way to incorporate incomplete information into commonly used macro models like the New Keynesian framework (Angeletos and Lian, 2018). Where my model differs from the existing literature is by studying a dynamic fundamental (i.e. the target variable for forecasts) and including a revised signal in the information set. This revised signal represents a new (noisy) observation of the fundamental value in the previous period, and thus captures the trade-off faced by the S.O.

Solving an incomplete information model with a dynamic fundamental is famously difficult, as Townsend (1983) pointed out this typically requires an infinite set of state variables. I take advantage of a recent methodological innovation by Huo and Pedroni (2020) who show the equilibrium action is equivalent to the optimal forecast of a modified problem that includes the degree of strategic incentives. Using their method I solve the first stage for individual actions given the precisions of all data signals and derive an expression for social welfare.

The second stage of the optimization problem has the S.O. maximize the social welfare function subject to their budget constraint by choosing two signal precisions: one for a current signal and one for a revised signal. The optimal choice of relative signal precisions reflects the marginal welfare trade-off between accuracy in the future versus timeliness today. The S.O. also considers the relative costs of increasing both signal precisions, with the assumption that increasing current precision is at least as expensive as increasing revised precision. This reflects the idea that marginal gains to accuracy prior to the initial publication date will involve additional expenses like overtime pay or high-wage temporary contractors.

I analyze the model solution under a variety of parameterizations and find the S.O.'s equilibrium choice can be a mixture both signal precisions, using the full budget on one signal instead of the other, choosing to not exhaust the budget or even providing neither signal. The mixture equilibrium results when agents have highly imprecise private information. In this case, the marginal gain to forecasting accuracy (and thus welfare) of increasing both signal precisions is high, and when the S.O. has a limited budget it chooses to spread resources across both current and revised signals. As private information becomes more precise the S.O. will increasingly shift resources from the revised signal into the current signal until it optimally chooses to only provide the latter, provided the degree of strategic incentives is small.

Large strategic incentives introduce the possibility of increased current precision crowding out private information. In a static beauty contest Morris and Shin (2002) (henceforth MS) showed how increasing the accuracy of information available to all agents can reduce social welfare when strategic incentives are high. The incentive to coordinate causes agents to excessively rely on public information at the expense of their private information because the former is useful both for forecasting the fundamental and predicting the average forecast. Since social welfare does not depend on the degree of dispersion in forecasts in a beauty contest, any weight placed on public information for coordination purposes is purely detrimental from a societal perspective.

I find a similar result to MS but with an important caveat. At a certain level of strategic incentives the S.O. will lower the precision of the public signal to avoid inefficient coordination, but they choose to redirect that budget back into the revised signal precision. This means welfare is decreasing in current precision but also increasing in revised precision at the threshold level of strategic incentives. The reason for this novel result is the difference in coordination value of both signals. Being a signal of the past, revised data are little use for agents trying to coordinate their forecasts, since any information about the past is heavily discounted through the forecasting mechanism. Despite this, the revised signal provides a net benefit to society because its small forecasting value outweighs its even smaller coordination value. This gives the S.O. a new lever to use when strategic incentives are high, whereby they can reduce inefficient coordination on public current information while still providing some useful public revised information.

When private information is very precise and strategic incentives are very high, the S.O. chooses to publish neither the current nor revised signal since either would risk crowding out the superior forecasting value of private information. In this parameter range the solution returns to that from MS, where any public information provision is welfare-decreasing.

The threshold levels of strategic incentives for which the revised signal is preferred or neither signal is offered depends on the size of the budget constraint and the relative costs. As the budget limit rises the S.O. puts more resources into the current signal, raising the level of strategic incentives before it switches back to the revised signal. This is due to a higher budget afford a more precise current signal which provides superior forecasting value compared to the revised signal, all else equal. Higher strategic incentives will eventually boost the costs of coordination above their forecasting value, but this threshold is pushed further out when there are more total resources available. In the limit with an unconstrained budget, the S.O. would choose a perfectly revealing current signal and no revised signal, just as in the idealized world.

My findings highlight the importance for the S.O. to understand three key features of their unique environment: i) the accuracy of private information; ii) the strength of strategic incentives; iii) the impact of shifting resources between current and revised signals. The first two are primitive parameters the S.O. takes as given in the model and in the real world, yet they determine which optimal choice set of signal precisions to offer. The last refers to the production function of the S.O. when choosing signal precisions. I assume linear marginal costs in the model, but in reality there are likely increasing resources required for similar improvements in current precision, as an example.

This paper contributes to the literature studying optimal data provision in the context of a beauty contest. As previously mentioned, MS began this research agenda by first showing how increased public information can reduce social welfare. Svensson (2006) argues that public information is welfare-enhancing under conservative parameterizations for private and public precisions within the set-up of MS. Other authors constructed different microfounded environments and found welfare weakly increases with public information, such as the investment game of Angeletos and Pavan (2004) or the monetary economy of Hellwig (2010). Angeletos and Pavan (2007) study a general set-up which includes MS as a special case, finding the MS result only holds due to a lack of other market frictions and coordination being socially undesirable. Amador and Weill (2012) show how the MS result can cause an informational externality whereby endogenously-generated signals become less informative of private information due to excessive coordination.

The literature listed above considers a static fundamental, so my paper is more closely related to recent work by Angeletos and Lian (2018) and Angeletos and Huo (2021) who study higher-order beliefs in a general equilibrium model with a dynamic fundamental. These papers include strategic incentives in the form of expectations of others' future actions, rather than their current actions. They also do not consider the question of optimal data provision. The methodology paper Huo and Pedroni (2020) applies their solution technique to a similar dynamic process as this paper but lacks the richer information framework I use that includes revised data as well as the optimization problem of the S.O.

Another related literature deals with modeling the decision problem of the S.O. directly, starting with Sargent (1989). He considers two possible models of the S.O., one in which the S.O. reports the true data measured with some noise or one in which the S.O. collects noisy data and filters it before publication. My approach is aligned with the second model, as it is the presence of noisy data that creates the trade-off between providing imprecise data now or more precise data next period. To my knowledge this is the first paper to study optimal data provision by the S.O. within the context of an incomplete information model.

2.2 Description of data revisions

Most macroeconomic data series have several different estimates published over time, with subsequent estimates reflecting revisions to the first publication. These data revisions are not corrections of previous errors, but instead reflect the receipt of additional source data that are used to produce superior estimates of the variable in question. The size of data revisions can be large: for example, the average revision to nominal GDP quarterly growth averaged 0.47 over a 40 year period (Aruoba, 2008). A recent extreme example occurred when real GDP growth in the first quarter of 2023 was initially estimated at 1.1%, but after revisions it is now estimated to be 2.2% (U.S. Bureau of Economic Analysis).

Statistical offices produce revised estimates in order to use more complete underlying source data. As an example, the first estimate of GDP produced by the Bureau of Economic Analysis is calculated by using four broad categories of data, ordered by how representative they are of the full population: comprehensive and near comprehensive, direct indicator, indirect indicator, and trend-based. As the name suggests, comprehensive data provide nearly complete coverage of the target population. The next two categories use extrapolation of related source data at a different frequency (e.g. using annual firm profits to estimate quarterly profits). The last category uses time series models to estimate source data when collected results are not available. Typically this is when survey collection is too slow to be incorporated into the first estimate (Holdren, 2014).

Many of the source data series that are collected from households involve labor-intensive survey methods, but as a result these tend to have higher response rates. The Current Population Survey (CPS), which provides data on labor force activities, is a questionnaire conducted in person or over the phone with roughly 60,000 households each month. It boasts the highest household response rate of 70% compared with, for example, the Consumer Expenditure (CE) survey which has a response rate of 41%. Surveys of establishments tend to use more automated tools like computer-assisted telephone interviews (CATI) or online forms, but these can suffer from low response rates. The Job Openings and Labor Turnover Survey (JOLTS), which is used for firm-side labor market statistics like job vacancies, has a response rate of just 33%, despite the collection period for the JOLTS being 4.5 times as long as for the CPS (U.S. Bureau of Labor Statistics).

Statistical offices are well-aware of the trade-off between high labor costs and receiving timely and complete responses. Regarding the JOLTS, the BLS requested an addition $9.6 million in funding (roughly 3% of the entire Labor Force Survey budget) in their 2024 budget justification in order to increase the sample of establishments and release the survey results two weeks earlier (U.S. Department of Labor, BLS-14). The same budget justification requests an additional $1 million to improve the timeliness of chained CPI measured by three months, as the current estimates are held up due to delays receiving current period expenditure weights from the CE survey (U.S. Department of Labor, BLS-46). When describing the need for additional funds, the BLS explained:

> The very features that make CPS unique and so valuable (high response rates and timely data) are the very same features that are proving problematic in the current environment. The extensive use of in-person field interviewers to collect data makes these two features possible. However, data collection costs are the most expensive component of the CPS program and are continuing to rise at an unsustainable pace. Mandatory pay and benefit increases for field representatives are outpacing appropriated funding and other costs (e.g., travel for personal visits) are increasing, while survey collection has become more challenging, requiring additional labor per case. (U.S. Department of Labor, BLS-29)

Given the rising costs associated with acquiring high quality and comprehensive data it is important to understand the value of additional data with respect to an objective measure. In the next section I describe the model set-up which relates the budget trade-offs described above to a metric of social welfare.

2.3 Model

Following MS I construct a simple economy which resembles a beauty contest over a single economic variable. The economic variable has persistence as in Huo and Pedroni (2023), but I add to their information structure an additional public signal of the past to represent data revisions. Agents use all sources of information to maximize their individual utility, which depends on their forecast accuracy and their distance from the average forecast. Finally I embed this beauty contest within the optimal data provision problem for the S.O., who chooses public signal precisions in order to maximize social welfare.

2.3.1 Information structure and individual behavior

The economic fundamental is a univariate persistent variable x_t which is never perfectly observed by anyone. Agents observe three noisy signals of the fundamental: a private signal that is unique to them and two public signals provided by the S.O. One public signal is called the current signal because it contains information about x_t directly. The other public signal is called the revised signal because it contains information only about x_{t-1}. Both public signals are contaminated by noise terms that are independent with each other but common across all agents.

$$\text{Fundamental: } x_t = \rho x_{t-1} + w_t$$

$$\text{Private signal: } y_{it} = x_t + \psi_{it}$$

$$\text{Current signal: } s_t = x_t + \varepsilon_t$$

$$\text{Revised signal: } r_t = x_{t-1} + \nu_t$$

The fundamental shock and all noise terms are i.i.d. and independent of each other and across time: $w_t \sim N(0, \tau_w^{-1})$, $\psi_{it} \sim N(0, \tau_\psi^{-1})$, $\varepsilon_t \sim N(0, \tau_\varepsilon^{-1})$, $\nu_t \sim N(0, \tau_\nu^{-1})$. All shocks are common across agents except for the variance of the private signal, τ_ψ^{-1}. I represent the variances of each shock as the inverse of τ which is the precision of the signal. For example, when the variance of the current signal goes to 0 this means the precision goes to infinity, and the current signal perfectly reveals the fundamental.

There is a continuum of agents indexed by i. Each agent uses the information structure to choose an action a_{it} to maximize the expectation of their period utility function:

$$u_{it}(a_{it}, a_t, x_t) = -\left[(1-\alpha)(a_{it} - x_t)^2 + \alpha(L_{it} - \bar{L}_t)\right] \tag{2.1}$$

where $a_t = \int a_{it} di$ is the average action across all agents. The first term represents the individual's incentive to accurately forecast the fundamental. The second term represents the beauty contest nature of the economy, where $L_{it} = \int_0^1 (a_{jt} - a_{it})^2 dj$ is the loss agent i suffers by choosing an action far away from the actions chosen by other agents. The weight agents assign to their forecast accuracy versus how close their action is to the average is governed by the parameter $\alpha \in [0, 1]$, which represents the strength of strategic incentives. All else equal, a larger strategic incentive means agents will choose actions they think are more alike at the cost of their individual forecast accuracy.

Taking the first-order condition in expectation gives agent i's optimal action:

$$a_{it} = (1-\alpha)\mathbb{E}_t[x_t|\Omega_{it}] + \alpha\mathbb{E}_t[\bar{a}_t|\Omega_{it}] \tag{2.2}$$

where Ω_{it} is the information set available to agent i in period t. Agents choose a_{it} to be a weighted average of their expectation of the fundamental x_t and their expectation of the average action, $\bar{a}_{it}$. The weight is α: when there are no strategic incentives ($\alpha = 0$) agents choose their best (mean-squared error minimizing) forecast. When $\alpha = 1$ agents want the same action as everyone else regardless of the fundamental.[1]

Period social welfare is the average across individual utilities:

$$\mathbb{W}_t(a_t, x_t) = \frac{1}{1-\alpha}\int u_{it}(a_{it}, a_t, x_t)di$$

where the term $\frac{1}{1-\alpha}$ normalizes social welfare so that it no longer depends on the degree of strategic incentives:

$$\mathbb{W}_t(a_t, x_t) = -\int (a_{it} - x_t)^2 di \tag{2.3}$$

Comparing equations (2.2) and (2.3) reveals the key tension in beauty contest economies. When $\alpha > 0$, individuals wish to minimize their forecast error as well as minimize their distance from the average action, which means their optimal choice will not minimize their forecast error. Social welfare, however, does not consider action dispersion and only cares about forecast accuracy, so any non-zero incentive to coordinate actions must reduce welfare. As a result, providing agents with more useful public signals can decrease welfare if agents use the signals to excessively coordinate their actions. In the static beauty contest of MS, increasing the precision of the public signal will reduce welfare if the ratio of public to private precisions is sufficiently small:[2]

$$\frac{\tau_\varepsilon}{\tau_\psi} \leq (2\alpha - 1)(1-\alpha) \tag{2.4}$$

When private information is relatively precise and $\alpha > 0.5$, increasing τ_ε decreases welfare due to agents over-relying on their public signal in order to choose actions that are similar to each other. In this sense public information crowds-out private information: private information improves welfare by improving forecast accuracy, but agents partially ignore this and choose actions that rely more on public signals due to their coordination incentive. In the next section I will show that including an additional public signal of revised data can overturn this result: it can be welfare enhancing to increase the precision of the revised signal due to its relatively low

[1] In this case, all agents choose the unconditional mean of x_t.

[2] In the notation from their paper, $\tau_\varepsilon = \alpha$, $\tau_\psi = \beta$, and $\alpha = r$.

coordination value compared with the current signal.

2.3.2 Optimal data provision

Given the beauty contest described above, the S.O. chooses the degree of precision for both the current (τ_ε) and revised signals (τ_v) in order to maximize expected social welfare. The S.O. is subject to a budget constraint which makes increasing signal precisions cost additional limited resources (e.g. labor hours). The total resources available to the S.O. are fixed by parameter I. For reasons described in section 2.2 I assume increasing the current signal precision is at least as costly as increasing the revised signal precision. The exact marginal cost of current signal precision is determined by the parameter $\phi \geq 1$. The S.O. solves

$$\begin{aligned} &\max_{\tau_\varepsilon, \tau_v} \ \mathbb{E}[\mathbb{W}(\tau_\varepsilon, \tau_v)] \\ &\text{s.t. } \ \phi\tau_\varepsilon + \tau_v \leq \text{I} \end{aligned} \tag{2.5}$$

where $\mathbb{W}$ is given by equation (2.3) and the expectation is taken with respect to the joint distribution of the fundamental shock and all noise terms. The model is solved in two steps:

1. Given τ_ε, τ_v, and parameters, agents choose a_{it} to maximize $u_{it}(a_{it}, a_t, x_t)$ conditional on their signals y_{it}, s_t, and r_t. Averaging across individual actions gives the aggregate action a_t.

2. Given a_t, the S.O. chooses τ_ε and τ_v to maximize expected period welfare subject to its budget constraint.

The S.O. views the following parameters as primitives for its optimization problem: properties of the fundamental ρ and τ_w; precision of private information τ_ψ; degree of strategic incentives α; and budget constraint parameters ϕ and I. It takes these as given, combined with the optimal actions of all agents, and chooses public signal precisions to maximize expected welfare.

2.4 Solution

I solve the model in two parts: finding equilibrium actions for each agent, and finding the optimal degree of public precisions chosen by the S.O.

2.4.1 Equilibrium action

Equation (2.2) provides each agent's best-response function: the optimal action given their information set and expectation of the average action across agents. Agents must solve their signal extraction problem to know $\mathbb{E}_t[x_t|\Omega_{it}]$ and also compute the functional form of $\mathbb{E}_t[\bar{a}_t|\Omega_{it}]$. In a static environment as in MS, this can be solved using guess-and-verify. The model of this paper involves a persistent fundamental which thus introduces the infinite regress problem identified by Townsend (1983). To see this, consider the following expansion of agent i's expectation of $\bar{a}_t$:[3]

$$\begin{aligned}
\mathbb{E}_{it}[\bar{a}_t] &= \mathbb{E}_{it}\left[\int_0^1 a_{it}di\right] = \mathbb{E}_{it}\left[(1-\alpha)\int_0^1 \mathbb{E}_{jt}[x_t]dj + \alpha\int_0^1 \mathbb{E}_{jt}[\bar{a}_t]dj\right] \\
&= (1-\alpha)\mathbb{E}_{it}\left[\bar{\mathbb{E}}x_t\right] + \alpha\mathbb{E}_{it}\left[\int_0^1 \mathbb{E}_{jt}\left[(1-\alpha)\int_0^1 \mathbb{E}_{kt}[x_t]dk + \alpha\int_0^1 \mathbb{E}_{kt}[\bar{a}_t]dk\right]dj\right] \\
&= (1-\alpha)\left(\mathbb{E}_{it}\bar{\mathbb{E}}x_t + \alpha\mathbb{E}_{it}\bar{\mathbb{E}}\bar{\mathbb{E}}x_t + \alpha^2\mathbb{E}_{it}\bar{\mathbb{E}}\bar{\mathbb{E}}\bar{\mathbb{E}}x_t + \dots\right) \\
&= (1-\alpha)\left(\mathbb{E}_{it}\bar{\mathbb{E}}x_t + \alpha\mathbb{E}_{it}\bar{\mathbb{E}}^2x_t + \alpha^2\mathbb{E}_{it}\bar{\mathbb{E}}^3x_t + \dots\right)
\end{aligned}$$

where the first line uses (2.2), the second line uses $\int_0^1 \mathbb{E}_{it}[x_t]di = \bar{\mathbb{E}}[x_t]$, and the fourth line collapses $\bar{\mathbb{E}}^2[x_t] = \bar{\mathbb{E}}\bar{\mathbb{E}}[x_t]$, $\bar{\mathbb{E}}^3[x_t] = \bar{\mathbb{E}}\bar{\mathbb{E}}\bar{\mathbb{E}}[x_t]$, and so on. With this expression of $\mathbb{E}_{it}[\bar{a}_{it}]$ in terms of the underlying fundamental, I can rewrite (2.2) as follows:

$$\begin{aligned}
a_{it} &= (1-\alpha)\mathbb{E}_{it}[x_t] + \alpha\mathbb{E}_{it}[\bar{a}_t] = (1-\alpha)\left(\mathbb{E}_{it}[x_t] + \alpha\mathbb{E}_{it}\bar{\mathbb{E}}[x_t] + \alpha^2\mathbb{E}_{it}\bar{\mathbb{E}}^2[x_t] + \dots\right) \\
&= (1-\alpha)\sum_{k=0}^{\infty}\alpha^k\mathbb{E}_{it}\left[\bar{\mathbb{E}}^k x_t\right]
\end{aligned}$$

[3] For readability I abbreviate $\mathbb{E}[x_t|\Omega_{it}] = \mathbb{E}_{it}[x_t]$.

This expression shows that the equilibrium action a_{it} depends on the agent's expectations of the fundamental plus a linear combination of all higher-order expectations, which are beliefs of others' beliefs. Since agents have different information sets through their idiosyncratic private signals, the law of iterated expectations does not apply, so the higher-order expectations do not collapse to agent i's first-order expectations. Agent i must therefore predict what everyone else will do, what everyone else thinks everyone else will do, and so on. Townsend (1983) showed computing all of these higher-order expectations requires infinite state variables to form priors for the cascading sequence of beliefs on beliefs on beliefs etc., which is computationally infeasible.

There are two main strategies for overcoming this problem. First is to convert the problem into the frequency domain and apply analytic techniques, as used in Kasa (2000) and Rondina (2008). This method provides neat closed-form forecasting equations that are useful for intuition, but generate messy welfare equations. The second method from Huo and Pedroni (2020) is to solve a forecasting problem after modifying the signal structure by the degree of complementarities. Although this provides a recursive form of the equilibrium action, I can study the the welfare expression numerically. To use this method, I first rewrite the information structure in state-space form:

$$\begin{aligned}
\boldsymbol{\xi}_{t+1} &= \boldsymbol{F}\boldsymbol{\xi_t} + \boldsymbol{u}_{t+1} \\
z_{it} &= \boldsymbol{H'}\boldsymbol{\xi_t} + \boldsymbol{\zeta_t} \\
\boldsymbol{\xi}_t &= \begin{bmatrix} x_t \\ x_{t-1} \end{bmatrix} \quad \boldsymbol{F} = \begin{bmatrix} \rho & 0 \\ 1 & 0 \end{bmatrix} \quad \boldsymbol{u_t} = \begin{bmatrix} w_t \\ 0 \end{bmatrix} \\
\boldsymbol{H'} &= \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \boldsymbol{\zeta_t} = \begin{bmatrix} \psi_{it} \\ \varepsilon_t \\ \nu_t \end{bmatrix}
\end{aligned}$$

Next I use the Kalman filter to derive optimal forecasts for the fundamental:

$$\begin{aligned} \hat{\boldsymbol{\xi}}_{t|t} &= \boldsymbol{F}\hat{\boldsymbol{\xi}}_{t-1|t-1} + \boldsymbol{K}(\boldsymbol{z}_{it} - \boldsymbol{H}'\boldsymbol{F}\hat{\boldsymbol{\xi}}_{t-1|t-1}) \\ \Rightarrow \hat{x}_{t|t} &= (\rho - \rho k_{11} - \rho k_{12} - k_{13})\hat{x}_{t-1|t-1} + k_{11}y_{it} + k_{12}s_t + k_{13}r_t \end{aligned} \tag{2.6}$$

where $\hat{x}_{t+j|t-s}$ is the forecast of x_{t+j} given information in time $t-s$. $\boldsymbol{K}$ is the steady-state Kalman gain matrix; the first row corresponds to the weight placed on the signals when forecasting the current fundamental x_t, and the second row contains the weights used for forecasting the fundamental last period x_{t-1}.[4] The Kalman gain matrix is a function of the squared forecast error matrix, $\boldsymbol{P}_{t+1} = \mathbb{E}[\hat{\boldsymbol{\xi}}_{t+1|t} - \boldsymbol{\xi}_{t+1}]^2$, which I find by solving the Ricatti equation:

$$\boldsymbol{P}_{t+1} = \boldsymbol{F}(\boldsymbol{P}_{t+1} - (\boldsymbol{P}_{t+1}\boldsymbol{H}(\boldsymbol{H}'\boldsymbol{P}_{t+1}\boldsymbol{H} + \boldsymbol{R})^{-1}\boldsymbol{H}'\boldsymbol{P}_{t+1}))\boldsymbol{F}' + \boldsymbol{Q} \tag{2.7}$$

where $\boldsymbol{Q} = \mathbb{E}[\boldsymbol{u_t}\boldsymbol{u_t'}]$ and $\boldsymbol{R}$ is the covariance matrix of the signals, $\mathbb{E}[\boldsymbol{\zeta_t}\boldsymbol{\zeta_t'}]$.

Applying the Kalman filter to this state-space framework would compute the optimal forecast for an individual agent given signals $\boldsymbol{z}_{it}$, but this is only part of their equilibrium action in (2.2). Importantly this ignores the coordination incentive (i.e. it computes a^*_{it} for $\alpha = 0$). Huo and Pedroni (2020) propose a method for finding the full form of (2.2) by solving a single agent's signal extraction problem after modifying their signals. The authors show that scaling the variance of all private noise terms by $\frac{1}{\sqrt{1-\alpha}}$ and solving for the optimal forecast obtains the same function of signals as the equilibrium action. In my framework this means modifying $\boldsymbol{\zeta_t}$ and $\boldsymbol{R}$ as follows:

$$\tilde{\boldsymbol{\zeta}}_{\boldsymbol{t}} = \begin{bmatrix} \frac{\psi_{it}}{\sqrt{1-\alpha}} \\ \varepsilon_t \\ \nu_t \end{bmatrix} \quad \tilde{\boldsymbol{R}} = \begin{bmatrix} (\tau_\psi(1-\alpha))^{-1} & 0 & 0 \\ 0 & \tau_\varepsilon^{-1} & 0 \\ 0 & 0 & \tau_\nu^{-1} \end{bmatrix}$$

[4] I assume $\rho < 1$ so the the fundamental is stationary and the Kalman gain matrix will converge to a steady-state value.

I denote variables determined by the alpha-modified signal as the original variable with a tilde over it, so $\boldsymbol{K}$ is the matrix of weights used purely for forecasting while $\tilde{\boldsymbol{K}}$ is the matrix of weights used in the equilibrium action. By replacing $\boldsymbol{R}$ in (2.7) with $\tilde{\boldsymbol{R}}$ and solving for $\tilde{\boldsymbol{P}}_{t+1}$ I find $\tilde{\boldsymbol{K}}$. Plugging this into equation (2.6) gives me the equilibrium action for agents as a function of their signals:

$$a_{it}^* = (\rho - \rho\tilde{k}_{11} - \rho\tilde{k}_{12} - \tilde{k}_{13})a_{it-1}^* + \tilde{k}_{11}y_{it} + \tilde{k}_{12}s_t + \tilde{k}_{13}r_t \tag{2.8}$$

where $\tilde{k}_{11}, \tilde{k}_{12}, \tilde{k}_{13}$ are all between $(0,1)$. The degree of strategic incentives, α, is embedded in each Kalman gain parameter. When $\alpha = 0$, $\tilde{k}_s = k_s$ for all s. Plugging equation (2.8) into (2.3) and solving gives the following expression for expected welfare:

$$\mathbb{E}[\mathbb{W}] = \frac{\tau_w^{-1}(1 - \tilde{k}_{11} - \tilde{k}_{12})^2 + \tau_\psi^{-1}\tilde{k}_{11}^2 + \tau_\varepsilon^{-1}\tilde{k}_{12}^2 + \tau_v^{-1}\tilde{k}_{13}^2}{1 - (\rho - \rho\tilde{k}_{11} - \rho\tilde{k}_{12} - \tilde{k}_{13})^2} \tag{2.9}$$

Welfare is a complicated function of the primitive parameters ρ, α, τ_ψ and the choice variables τ_ε, τ_v, since they all have non-linear relationships with each modified Kalman gain parameter.

2.4.2 Optimal data provision

The solution for S.O.'s problem is straightforward in abstract terms. Taking first-order conditions and re-arranging gives the following two conditions dictating the choice of τ_v^* and τ_ε^*:

$$\frac{\frac{\partial \mathbb{W}}{\partial \tau_\varepsilon}}{\frac{\partial \mathbb{W}}{\partial \tau_v}} = \phi$$

$$\phi\tau_\varepsilon + \tau_v = I$$

Using the specific expression for welfare from equation (2.9) quickly makes the problem analytically intractable, and it is not feasible to derive a closed-form solution for τ_v^* and τ_ε^*. Fortunately the expressions for welfare and its constituents (e.g. Kalman gains and mean-square error matrix

$\tilde{\boldsymbol{P}}_{t+1}$) are smooth and can be solved numerically. To proceed I solve for τ_ν^* and τ_ε^* using the equation solver in MATLAB and study comparative statics numerically.

2.5 Results

Numerically solving the problem requires choosing fixed parameter values. I set $\rho = 0.95$ to match highly-persistent data series as well as to best showcase the impact of revised data. All signal precisions are normalized by the inverse variance of the fundamental, which means $\tau_w = 1$ and signal precisions can be interpreted as signal-to-noise ratios.

I start by analyzing a benchmark set of parameter values which demonstrate the unique properties of revised information. Then I consider how changing the budget constraint affects the trade-offs between current and revised signal precisions.

2.5.1 Benchmark case - low τ_ψ

For the benchmark case I set the budget limit $I = 1$ and the cost of improving the current signal $\phi = 5$, so that the maximum attainable value of the current signal precision is $\tau_\varepsilon = 0.2$ and the marginal rate of technical substitution between the current and revised signal is $\frac{1}{5}$. Finally, I allow τ_ψ and α to vary and study the response of τ_ν^* and τ_ε^*.

Figure 2.1 shows the optimal choice of current and revised signal precisions for a selection of relatively small τ_ψ values and across the entire range of α. When individuals have very inaccurate private signals as in the top row, $\tau_\psi \in (0.1, 0.2)$, the welfare-maximizing choice of public signals puts more weight on the revised signal instead of the current signal. Both signals receive some weight and the budget constraint is fully exhausted, which is true for the entire range of α. The bottom row of the figure shows τ_ν^* and τ_ε^* when individuals have more accurate private signals, $\tau_\psi \in (0.3, 0.4)$. As τ_ψ rises the weight placed on the revised signal falls and the weight placed on the current signal rises, culminating in a range of α (roughly between 0 and 0.65) over which $\tau_\varepsilon^* > \tau_\nu^*$ when $\tau_\psi = 0.4$.

The overall lesson from figure 2.1 is that when individuals have poor private information

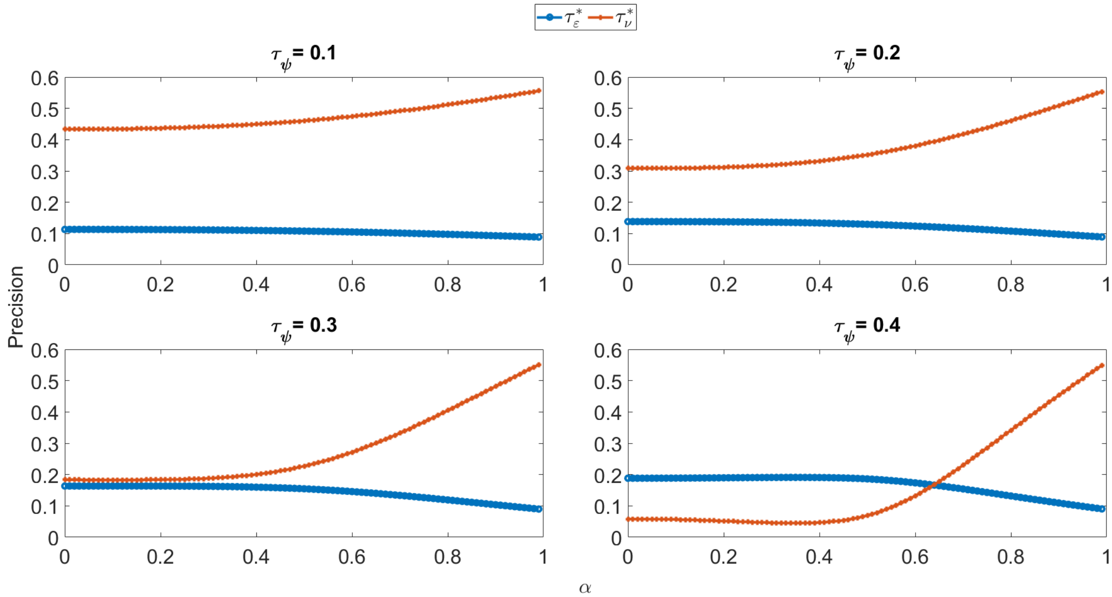

Figure 2.1. *Notes:* Each quadrant plots the optimal choices of τ_ν and τ_ε as α changes for a given level of τ_ψ. The blue line dotted with circles represents τ_ε^* and the red line dotted with pluses represents τ_ν^*. The figure uses the following fixed parameter values: $\rho = 0.95$, $\tau_w = 1$, $\phi = 5$, $I = 1$.

it is welfare-enhancing to provide them a mixture of signals about both the present and the past. The northwest plot shows that $\tau_\varepsilon^* \approx 0.1$ for the entire range of α, which amounts to only 50% of the available budget. The rest should be spent on improving the revised signal precision. A mixture of both signals is optimal even as individuals receive increasingly accurate private signals, as can be seen in the southeast plot which suggests that even the lowest value of τ_ν^* should be about 5% of the total budget.

The S.O. chooses higher precision for the revised signal than the current signal for two reasons: a small budget constraint and low private precision. First consider the impact of the budget constraint. If the S.O. was unconstrained they would set $\tau_\varepsilon = \infty$ and perfectly reveal the fundamental, rendering the revised signal irrelevant and maximizing welfare regardless the degree of strategic complementarities. When the budget constraint is small (i.e. I is small relative to τ_w) the best the S.O. could do using the current signal alone is set $\tau_\varepsilon = \frac{I}{\phi}$, which would be $\tau_\varepsilon = 0.2$ in the benchmark case. This does not maximize welfare because the marginal gains to forecasting accuracy of raising τ_v away from zero are large. Specifically, these gains are larger than the gains in forecast accuracy of increasing τ_ε from 0.1 to 0.2.

To demonstrate this I plot the mean-squared error of both signals as their precision rises in figure 2.2. This captures the pure forecasting returns to each signal (i.e. assuming $\alpha = 0$) since the findings in figure 2.1 are true when there are no strategic incentives.[5] Holding $\tau_\psi = 0.1$ and $\tau_v = 0$, the black dashed curve represents the forecast error for any choice of $\tau_\varepsilon \in (0.1, 0.2)$. This is a feasible choice set for the S.O. as outlined previously, but the slope of the curve is fairly flat, indicating marginal increases in current precision provide small and nearly constant reductions in errors. The red dotted curve represents how increasing revised signal precision reduces forecast error assuming $\tau_\varepsilon = 0.1$, which captures the trade-off inherent in figure 2.1. The steeper τ_v curve is why the S.O. chooses to put some of its budget in both signals, since the returns to the revised signal are large when $\tau_v = 0$ and $\tau_\varepsilon = 0.1$.[6]

Now consider the impact of the degree of private precision. When τ_ψ is small relative to τ_w, agents are poorly informed and thus rely heavily on public signals to improve their forecast accuracy. This means the reduction in MSE will be greater when increasing either τ_ε or τ_v if τ_ψ is small rather than when it is large. I show this in figure 2.3, which plots MSE when using either public signal on its own alongside different levels of private precision. As τ_ψ rises the slopes of both the current and revised signal curves flatten, meaning the largest gains to welfare from increasing public precisions comes when τ_ψ is small. This fact combined with the small budget constraint is what creates the favorable marginal reduction in MSE from having $\tau_v > \tau_\varepsilon$ as seen in figure 2.2.

The mechanism described thus far ignores the role of strategic incentives. These become more important as both τ_ψ and α rises, such as in the bottom row of figure 2.1.[7] For example,

[5] Strategic incentives play a vital role for other results from figure 2.1 which I discuss later.

[6] It is important to note that the shape of both curves crucially depends on the value of the other signal's precision. For a larger value of τ_ε the MSE curve when raising τ_v (i.e. the red dashed line) becomes very flat, which will explain results in the next section regarding lifting the budget limit.

[7] The graphs in the top row also demonstrate the impact of strategic incentives through the slight upward slope to the τ_v^* curve, but this becomes much more pronounced as τ_ψ increases.

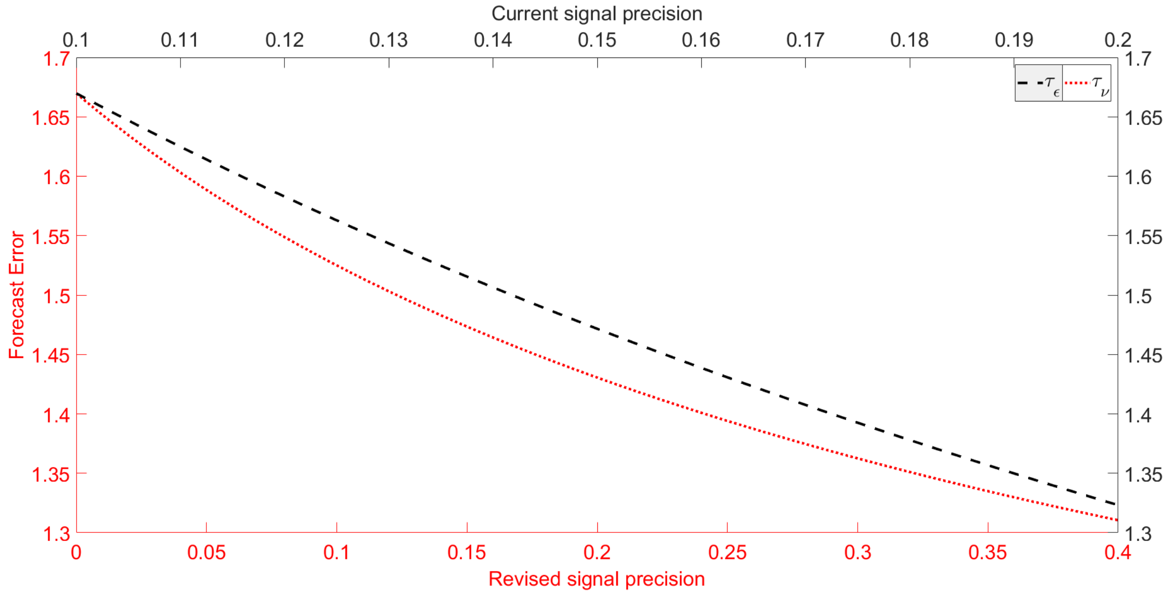

Figure 2.2. Forecast Error of Current and Revised Signals. The graph plots the mean-squared error for forecasts made under the following parameterizations: $\rho = 0.95, \tau_w = 1, \tau_\psi = 0.1, \alpha = 0$. The dashed black line represents the MSE when increasing τ_ε from 0.1 to 0.2, holding $\tau_\nu = 0$. The dotted red line represents the MSE when increasing τ_ν from 0 to 0.4, holding $\tau_\varepsilon = 0.1$.

the southwest plot with $\tau_\psi = 0.3$ shows the line for τ_ε^* shifting up and the line for τ_ν^* necessarily shifting down, as the S.O. exchanges revised signal precision for current signal precision. The two signal precisions are close together while strategic incentives are small (i.e. $\alpha < 0.4$) but begin to diverge significantly afterwards. As α approaches 1 the optimal choice of precisions returns to the case when private information was inaccurate.

More accurate private information increases the coordination costs of public information, as increasingly precise public signals will effectively crowd-out useful private signals. When the incentive to coordinate for individuals rises, agents will place excessive weight on public signals relative to their value from a forecasting perspective, ultimately causing a reduction in social welfare. This is the conclusion from MS, though it is largely a function of the beauty contest nature of the economy they study (Angeletos and Pavan, 2007). The economy modeled here is also a beauty contest but we see the opposite result due to the revised signal: increasing revised signal precision is welfare-enhancing when strategic incentives are very high. This increased

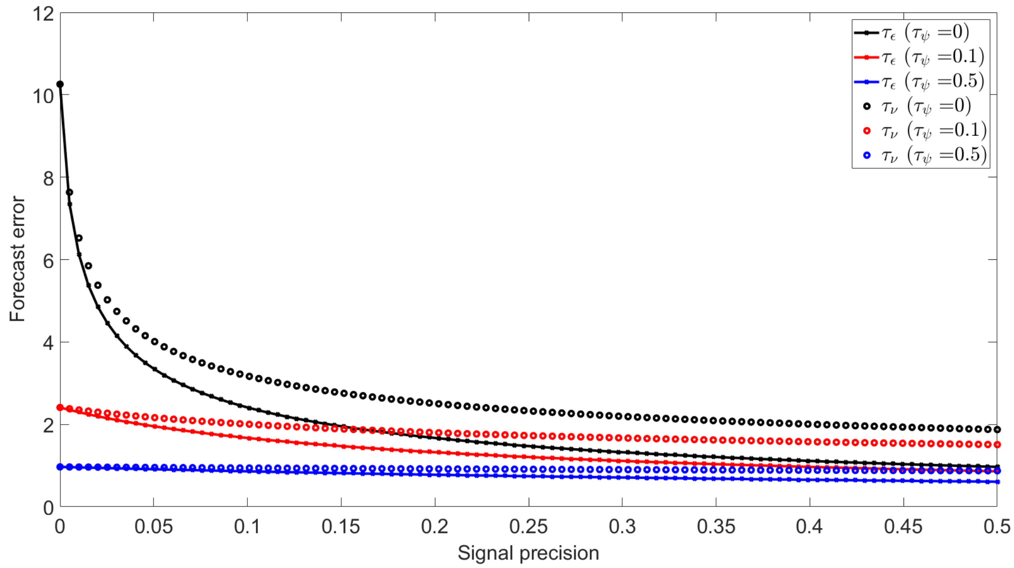

Figure 2.3. Forecast Error of Current and Revised Signals. The graph plots the mean-squared error for forecasts made under the following parameterizations: $\rho = 0.95, \tau_w = 1, \alpha = 0$. The solid lines with dots represent MSE when increasing τ_ε for various levels of τ_ψ while holding $\tau_\nu = 0$. The lines of circles represent MSE when increasing τ_ν for various levels of τ_ψ while holding $\tau_\varepsilon = 0$.

precision comes at the cost of decreased current signal precision, as the S.O. optimally balances the forecasting gains and coordination costs from the two signals.

A more accurate revised signal reduces the weight agents place on the current signal due to the mechanics of their forecasting problem. When solved with the Kalman filter, the weight on each signal corresponds to its associated Kalman gain (modified according to the procedure from section 2.4). Figure 2.4 plots the α-modified Kalman gain on the current signal as τ_ν changes, holding τ_ε equal to its optimal choice given similar parameters to figure 2.1. The slope of the curve represents how much agents will shift weight from the current signal to the revised signal as τ_ν increases. The downward slope is as expected - when the revised signal becomes relatively more informative, agents will put more weight on it at the expense of weight on the current signal.

Comparing across the curves gives a sense of the cross partial derivative with respect to τ_ν and α. Since the red dashed line is steeper than the blue crossed line, this cross-partial is

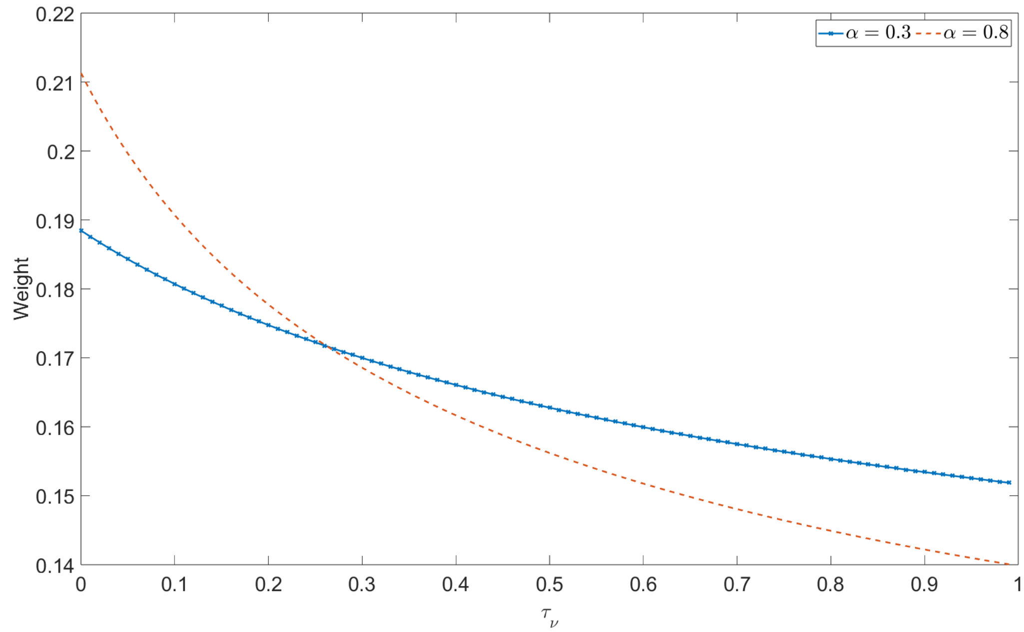

Figure 2.4. Weight placed on current signal. The degree of current precision is calculated as the optimal choice given the following parameter values: $\rho = 0.95$, $\tau_w = 1$, $\tau_\psi = 0.3$, and $\tau_\nu = \tau_\nu^*$. When $\alpha = 0.3$ as in the blue crossed line, $\tau_\nu^* = 0.188$. When $\alpha = 0.8$ as in the red dashed line, $\tau_\nu^* = 0.405$.

positive, meaning higher strategic incentives increase the impact of τ_ν on the current signal's weight. This explains the steep upward slope of the τ_ν^* line in the bottom row of figure 2.1: when α is large, the S.O. can minimize the coordination incentive for agents by increasing τ_ν and thus forcing agents to shift weight from the current signal to the revised signal. Agents will rationally adjust their weights because the revised signal provides increased forecasting accuracy.

In effect, the revised signal gives the S.O. a new lever for addressing socially inefficient coordination. Without a revised signal the only way to reduce excessive coordination is to lower current signal precision, which also lowers welfare through decreased forecast accuracy. With a revised signal the S.O. can both reduce excessive coordination and increase forecast accuracy by raising the revised signal precision.

The southeast quadrant of figure 2.1 shows how sufficiently accurate private information causes $\tau_\varepsilon^* > \tau_\nu^*$ when strategic incentives are small. This is due to shifting of the τ_ε^* and τ_ν^* curves as described above. The upward slope of the τ_ν^* curve becomes steeper for high α, in line with the mechanics underlying figure 2.4. The next section pushes τ_ψ further until the optimal choice of τ_ν^* is zero.

2.5.2 Benchmark case - medium τ_ψ

Figure 2.5 plots optimal precisions of the current and revised signal for significantly higher degrees of private precision. The graphs continue the pattern observed in figure 2.1 but now it is no longer always optimal to have a mixture of current and revised signal precisions. Each graph in the figure shows that $\tau_\nu^* = 0$ as long as α is below a specific threshold level, $\bar{\alpha}$. This threshold increases as the degree of private precision increases. Since $\tau_\nu > 0$ in each graph from figure 2.1, very small τ_ψ means $\bar{\alpha} = 0$ (even if $\tau_\varepsilon^* > \tau_\nu^*$ as in the southeast quadrant).

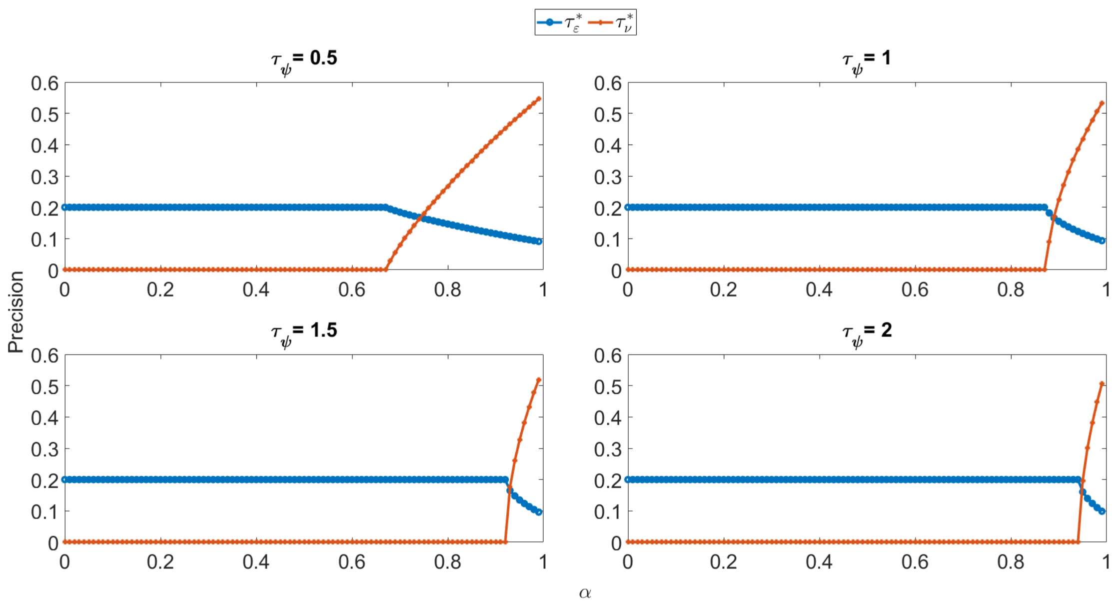

Figure 2.5. Optimal choice of τ_ν, τ_ε - benchmark parameters, medium τ_ψ. Each quadrant plots the optimal choices of τ_ν and τ_ε as α changes for a given level of τ_ψ. The blue line dotted with circles represents τ_ε^* and the red line dotted with pluses represents τ_ν^*. The figure uses the following fixed parameter values: $\rho = 0.95$, $\tau_w = 1$, $\phi = 5$, $I = 1$.

Holding α constant, a more precise private signal reduces the potential welfare loss due to public signal coordination. This is because higher τ_ψ will naturally reduce the weight agents place on their public information, similar to the mechanism in figure 2.4. Figure 2.6 repeats this exercise but instead holds τ_v fixed at its optimal choice given parameters and plots the weight on the current signal as τ_ψ increases. Again there is a downward slope when α is held constant, indicating that agents will optimally shift weight from the current signal to their private information as it becomes more accurate. This reduces the potential crowding-out of private information and brings private actions closer to the social-maximizing ones. When agents act accordingly the S.O. can put all its resources into improving the current signal which improves forecast accuracy by more than the revised signal.

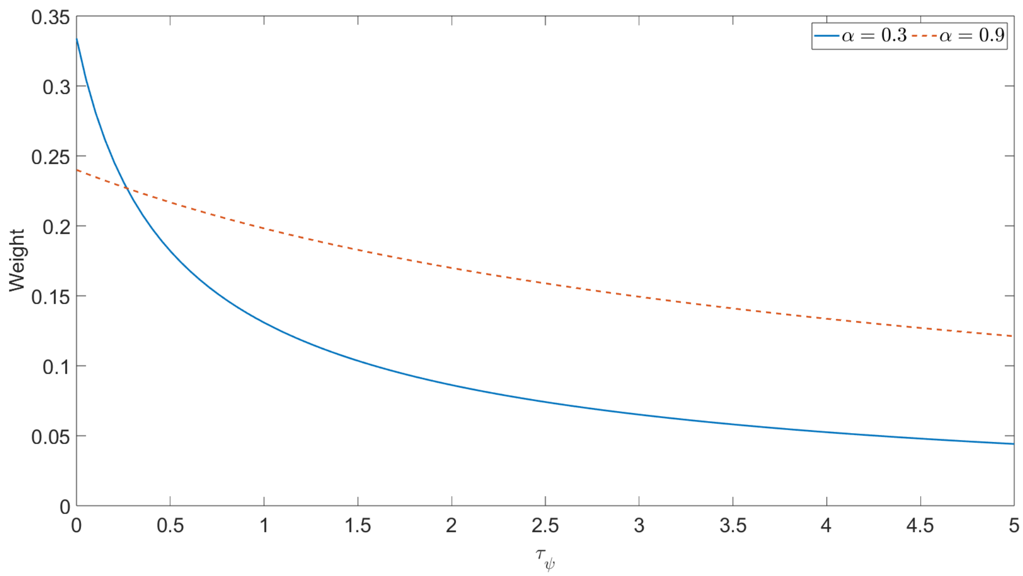

Figure 2.6. Weight placed on current signal. The degree of current and revised signal precisions is fixed at the optimal choice given the following parameter values: $\rho = 0.95$, $\tau_w = 1$. When $\alpha = 0.3$ as for the blue solid line, $(\tau_\varepsilon^*, \tau_v^*) = (0.2, 0)$. When $\alpha = 0.9$ as for the red dashed line, $(\tau_\varepsilon^*, \tau_v^*) = (0.155, 0.225)$. The optimal signal precisions $(\tau_\varepsilon^*, \tau_v^*)$ were computed assuming $\tau_\psi = 1$. Choosing a different τ_ψ for this computation does not materially change the graph.

However, this crucially relies on the level of α being below the threshold. Above this threshold agents will put too much additional weight on their public information at the expense of fully utilizing their accurate private information. This is evident in the far flatter slope of the red dashed line in figure 2.6: with $\alpha = 0.9$ agents have strong incentives to coordinate, thus they are less willing to shift weight from public information to their private signals even as τ_ψ rises. This is where the S.O. can improve welfare by both reducing the quality of the current signal and increasing the quality of the revised signal. Reducing τ_ε reduces the weight individuals place on the current signal for both forecasting and coordination purposes. Increasing τ_ν has two effects: it improves forecasting accuracy by more than the costs of coordination (since τ_ν was initially 0), and it shifts weight from τ_ε as described by figure 2.4. The mechanism is the same as when τ_ψ was low as in figure 2.1, but now there is a clear cut-off point of $\bar{\alpha}$ that informs the S.O. of whether it should put all its resources in the current signal ($\alpha < \bar{\alpha}$) or if it should mix its resources across current and revised signals ($\alpha > \bar{\alpha}$).

2.5.3 Benchmark case - high τ_ψ

Finally, I plot the benchmark case parameterization with relatively high values of τ_ψ in figure 2.7. The pattern from figure 2.5 continues in the northwest quadrant when $\tau_\psi = 3$, however once $\tau_\psi \geq 5$ there is a new equilibrium: no public information is provided (i.e. $\tau_\varepsilon^* = \tau_\nu^* = 0$). When private information is very precise, the S.O. chooses to provide no useful public information for high levels of strategic incentives in order to avoid crowding-out private information. This is analogous to the condition from MS which I reproduced in equation (2.4). The threshold level of strategic incentives at which public information becomes welfare decreasing, $\underline{\alpha}$, is decreasing in τ_ψ until it reaches a limiting value.

The S.O. chooses to not fully use the budget when strategic incentives and private precision are both very high. Figure 2.8 plots the surface of welfare for various choices of τ_ε and τ_ν, with the budget constraint represented by the red plane. Points on the welfare surface anterior to the red plane are infeasible, while points posterior to the red plane are feasible but do not exhaust the budget. Welfare slopes upwards away from the red plane as both τ_ε and τ_ν decline, reaching a maximum when they are both equal to 0. This is because increasing either signal precision leads to inefficient forecasting weights due to excessive coordination. If the S.O. was required to fully use their budget, they would put all their resources into τ_ε since the welfare surface is sloping downward away from this point at the back-left of the figure.

Next I study how the budget limit I and the marginal cost of current signal precision ϕ

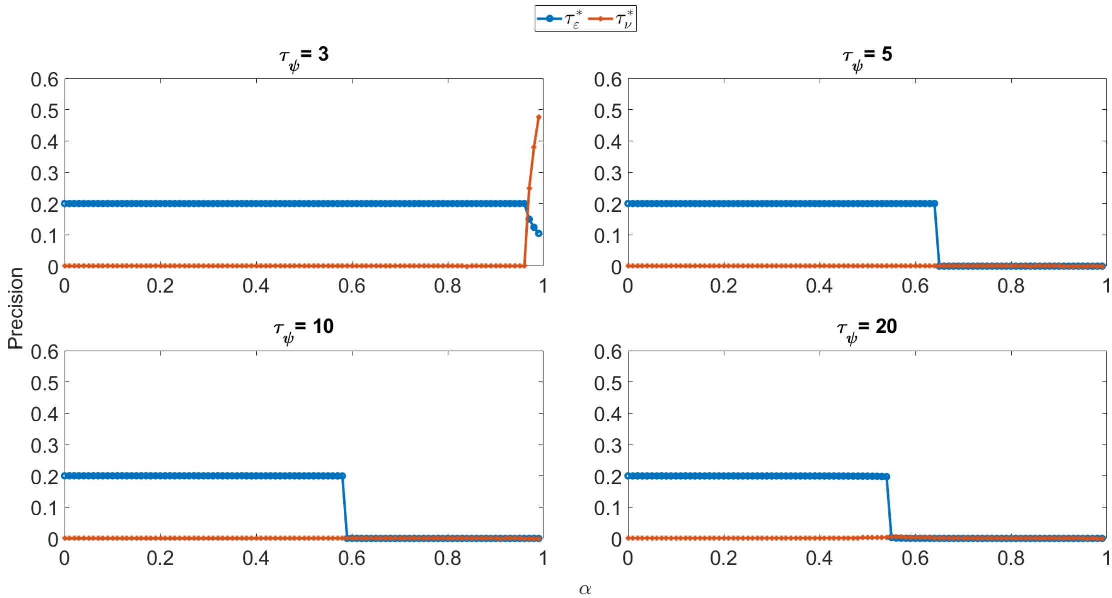

Figure 2.7. Optimal choice of τ_ν, τ_ε - benchmark parameters, high τ_ψ. Each quadrant plots the optimal choices of τ_ν and τ_ε as α changes for a given level of τ_ψ. The blue line dotted with circles represents τ_ε^* and the red line dotted with pluses represents τ_ν^*. The figure uses the following fixed parameter values: $\rho = 0.95$, $\tau_w = 1$, $\phi = 5$, $I = 1$.

affect this threshold $\bar{\alpha}$ and the trade-off between τ_ν^* and τ_ε^*.

2.5.4 Raising the budget limit

Figure 2.9 shows how the optimal choice of τ_ν and τ_ε changes as the budget limit doubles from 1 to 2. The most striking difference from the analogous figure 2.1 is that $\tau_\varepsilon^* > \tau_\nu^*$ for all levels of τ_ψ considered and over the entire range of α. The two curves have similar shapes as they did in figure 2.1, but the rise in the budget limit caused a large shift in opposite directions so that it is always optimal to put more resources into the current signal rather than the revised signal, even when strategic incentives are very high. Additionally, the threshold level of α at which all resources should be put into the current signal has increased to roughly $\bar{\alpha} = 0.52$ when $\tau_\psi = 0.3$. When the budget limit was lower, $\bar{\alpha} = 0$ until $\tau_\psi \geq 0.5$.

A higher budget limit lets the S.O. increase welfare by raising the accuracy of the current signal to a level that was previously unobtainable. Interestingly, τ_ε^* has more than doubled from its value under the benchmark parameters, from 0.1 to 0.4, suggesting there is a non-linear

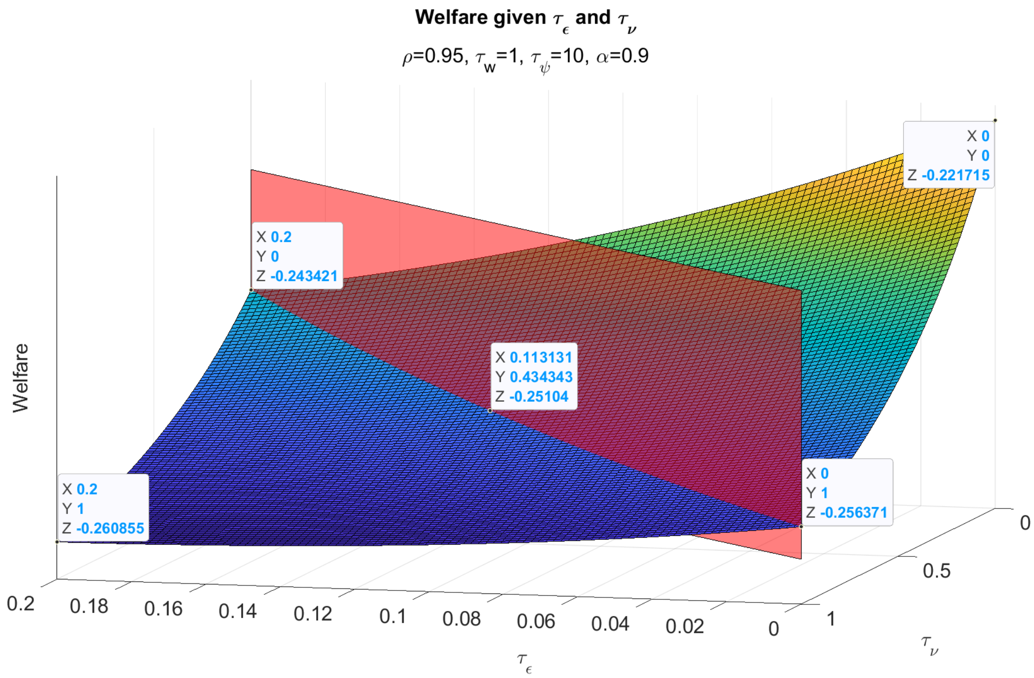

Figure 2.8. Welfare, high α and τ_ψ. The vertical axis (z) represents welfare given the choice of τ_ε (x-axis) and τ_ν (y-axis). The figure uses the following fixed parameter values: $\rho = 0.95$, $\tau_w = 1$, $\tau_\psi = 10$, $\alpha = 0.9$, and $\phi = 5$.

relationship between the budget limit and τ_ε^*. The new budget limit is still fully exhausted for all levels of τ_ψ, but only uses a mixture of current and revised signal precisions when τ_ψ is small.

In a sense, the S.O. views the revised signal as an inferior good (for a given level of α). Increasing the budget constraint allows the choices of both τ_ε and τ_ν to rise, but instead the S.O. chooses to decrease τ_ν^* in favor of increasing τ_ε^*. This is evident from figure 2.10 which plots the surface of τ_ν^* as the budget limit and strategic incentives change. Revised signal precision increases in the budget limit only over a very small range, since with such a low budget and high ϕ the forecasting gains from the current signal are too low to invest in. Around $I = 0.6$, the revised signal becomes an inferior good: increasing the budget limit leads to lower values of τ_ν^*. The reason is the relatively low forecasting value of the revised signal, which is clear when considering limits: a perfect current signal will eliminate forecast errors, but a perfect revised signal cannot.

The higher threshold level of α before the S.O. chooses to offer the revised signal again

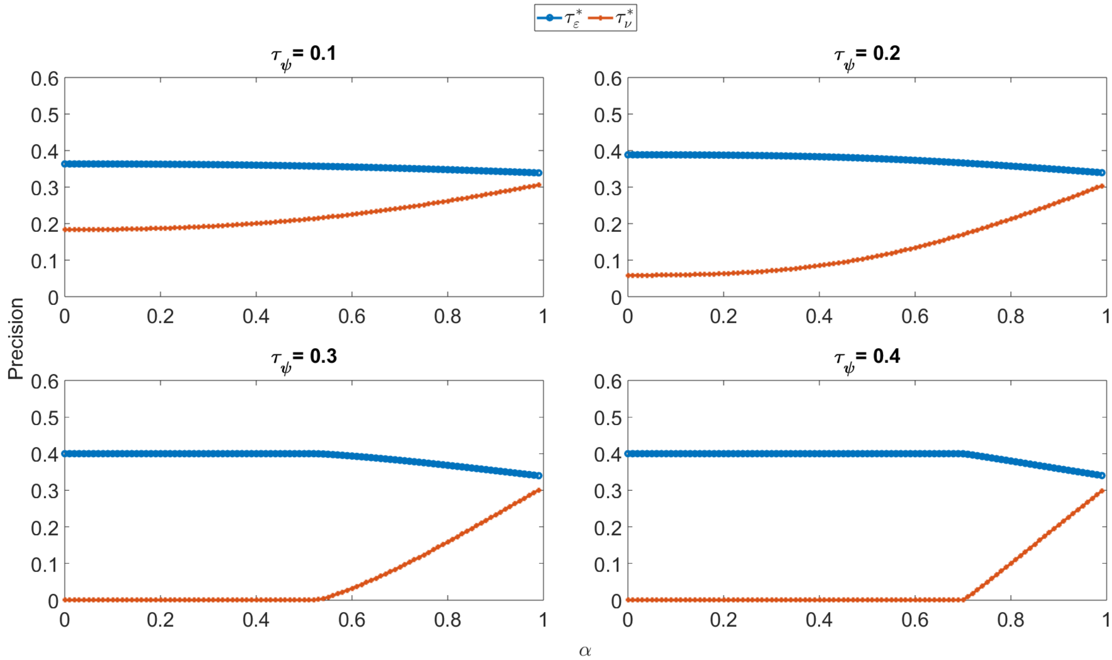

Figure 2.9. Optimal choice of τ_ν, τ_ε - budget limit = 2, low τ_ψ. Each quadrant plots the optimal choices of τ_ν and τ_ε as α changes for a given level of τ_ψ. The blue line dotted with circles represents τ_ε^* and the red line dotted with pluses represents τ_ν^*. The figure uses the following fixed parameter values: $\rho = 0.95$, $\tau_w = 1$, $\phi = 5$, $I = 2$.

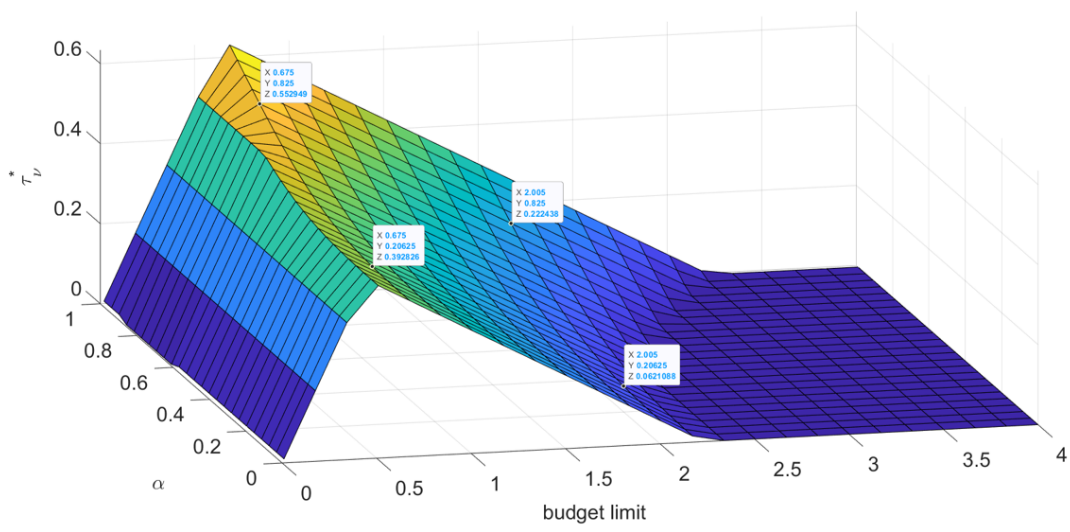

Figure 2.10. Optimal choice of τ_ν. The vertical axis (z) represents the optimal choice of τ_ν for a given α (y-axis) and budget limit (x-axis). The figure uses the following fixed parameter values: $\rho = 0.95$, $\tau_w = 1$, $\tau_\psi = 0.2$, $\phi = 5$.

is most evident when considering the high range of τ_ψ as in figure 2.11. Compared with the benchmark case, $\bar{\alpha}$ is significantly higher for all levels of τ_ψ, so the S.O. will only use the revised signal when the degree of strategic incentives is very high. This is because with a higher budget limit, welfare is enhanced by the greater forecast accuracy from putting more resources into the current signal. This pushes $\bar{\alpha}$ higher due to the improved accuracy overpowering the costs of coordination for a larger range of α.

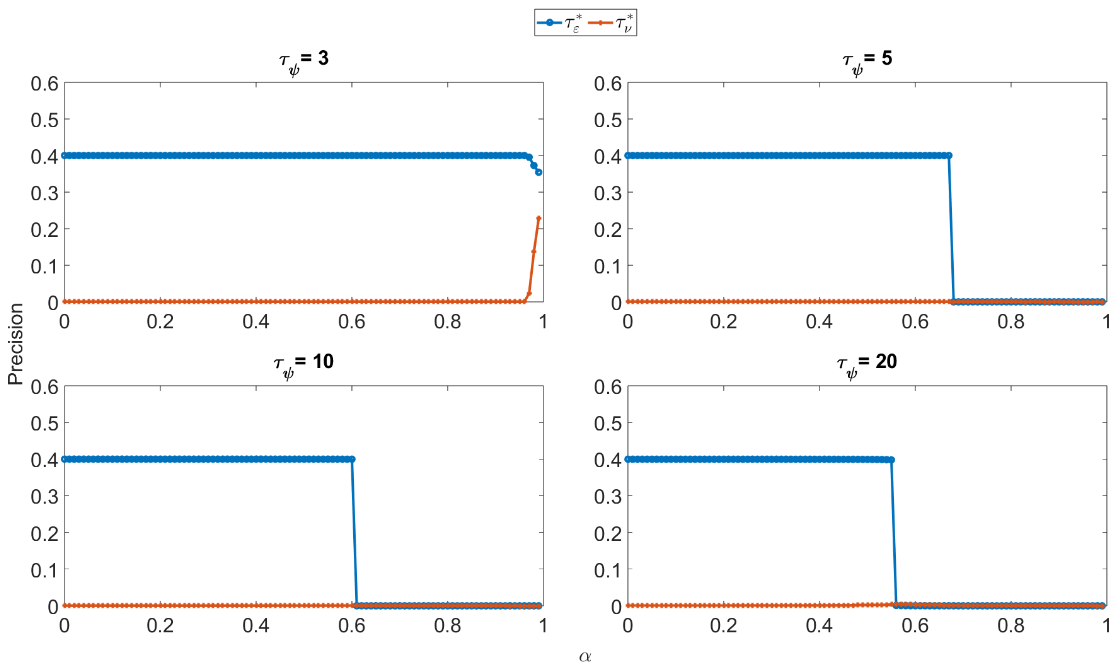

Figure 2.11. Optimal choice of τ_ν, τ_ε - budget limit = 2, high τ_ψ. Each quadrant plots the optimal choices of τ_ν and τ_ε as α changes for a given level of τ_ψ. The blue line dotted with circles represents τ_ε^* and the red line dotted with pluses represents τ_ν^*. The figure uses the following fixed parameter values: $\rho = 0.95$, $\tau_w = 1$, $\phi = 5$, $I = 2$.

To show how the budget limit affects the tension between forecasting gains and coordination costs, I plot the weights placed by agents on the current signal for both purposes in figure 2.12. The black solid and dashed lines correspond to the weights agents choose when $\tau_\varepsilon = 0.2$ as in the benchmark case studied in figure 2.7. The lower forecasting weight when τ_ε is small explains why $\bar{\alpha}$ is lower: the ratio between forecasting gains and coordination losses is smaller, so the S.O. begins mixing signal precisions for a lower level of α. When the budget limit rises and $\tau_\varepsilon = 0.4$ becomes obtainable, agents place more weight on the current signal due to its high forecasting gains. The coordination incentive changes little as τ_ε rises which means the S.O. will prefer providing only the current signal for a higher degree of α, thus increasing $\bar{\alpha}$.

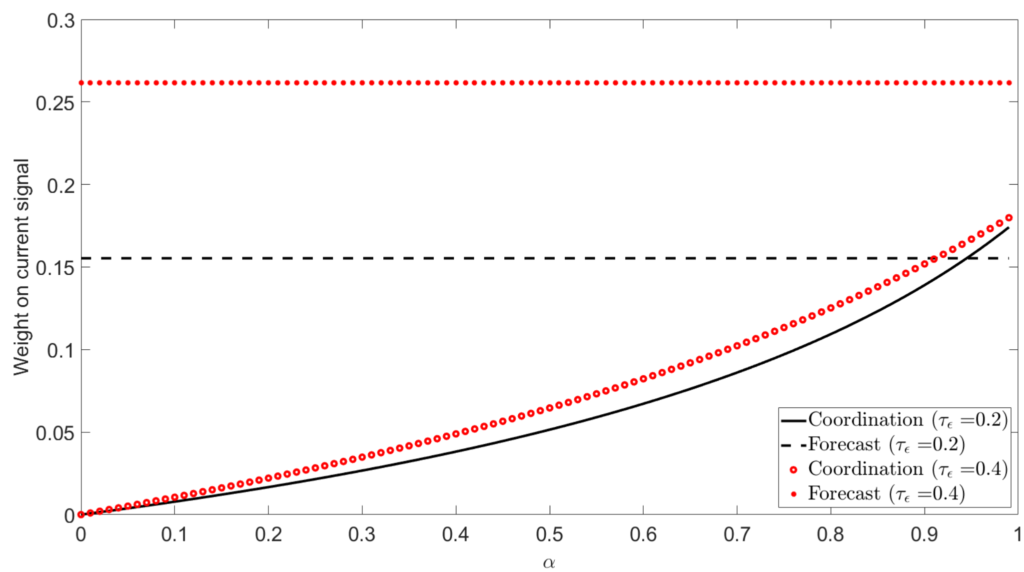

Figure 2.12. Weight placed on current signal - forecasting vs. coordination. Coordination weight (solid black and red circles) is computed as the total weight placed by agents on the current signal minus the weight they would have placed if $\alpha = 0$. Forecasting weight (dashed black and red dots) is computed as the total weight minus coordination weight. The figure uses the following fixed parameter values: $\rho = 0.95$, $\tau_w = 1$, $\tau_\psi = 0.5$, and $\tau_v = 0$.

As the budget limit continues to rise, the revised signal becomes increasingly irrelevant for all ranges of α and τ_ψ due to the vast improvements in forecasting accuracy afforded by the current signal. This effect is symmetric: when the budget limit decreases, the revised signal becomes increasingly more important and $\bar{\alpha}$ falls. Next I hold the budget limit fixed and allow the marginal cost of current signal precision, ϕ, to change.

2.5.5 Changing the cost of current signal precision

The size of the relative costs between both signal precisions has a similar relationship with τ_ε^* and τ_v^* as the budget limit. Lowering ϕ shifts τ_ε^* up and τ_v^* down in a manner reminiscent of figure 2.11. This is no surprise considering the mechanism underpinning that figure. When the budget limit rises, the S.O. can put more resources into the current signal which has higher forecasting value. A decrease in the marginal cost of the current signal also allows the S.O. to put more resources into improving its precision, so the effect is the same.

Raising ϕ has a similar impact as lowering the budget limit, since it makes the more valuable signal for forecasting less affordable. Figures 2.13 and 2.14 plot τ_ε^* and τ_ν^* when the cost of improving the current signal is 7.5 times that of the cost improving the revised signal (a 50% increase from the benchmark case). As expected, the revised signal has increased importance over a larger range of τ_ψ and even is the only signal provided when private precision is very small and α large. The threshold value $\bar{\alpha}$ decreased significantly from the benchmark; previously $\bar{\alpha} = 0.65$ when $\tau_\psi = 0.5$, but now with more expensive current signal precision $\bar{\alpha} = 0$ at the same degree of private precision. A more expensive current signal means it is harder for the S.O. to overcome the coordination costs, so it will increase τ_ν for lower levels of α.

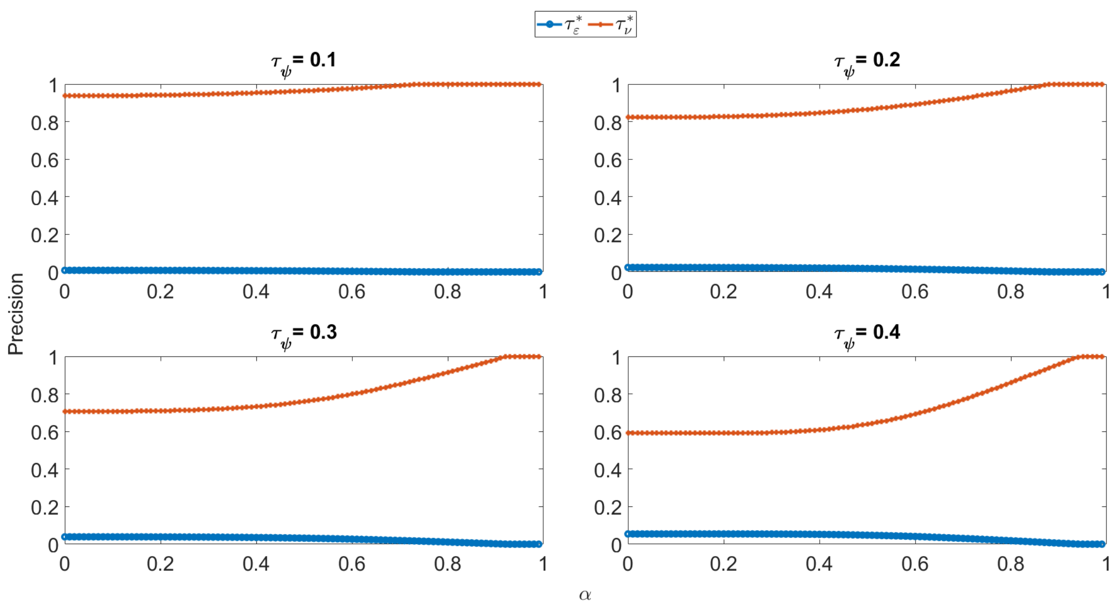

Figure 2.13. Optimal choice of τ_ν, τ_ε - $\phi = 7.5$, low τ_ψ. Each quadrant plots the optimal choices of τ_ν and τ_ε as α changes for a given level of τ_ψ. The blue line dotted with circles represents τ_ε^* and the red line dotted with pluses represents τ_ν^*. The figure uses the following fixed parameter values: $\rho = 0.95$, $\tau_w = 1$, $\phi = 7.5$, $I = 1$.

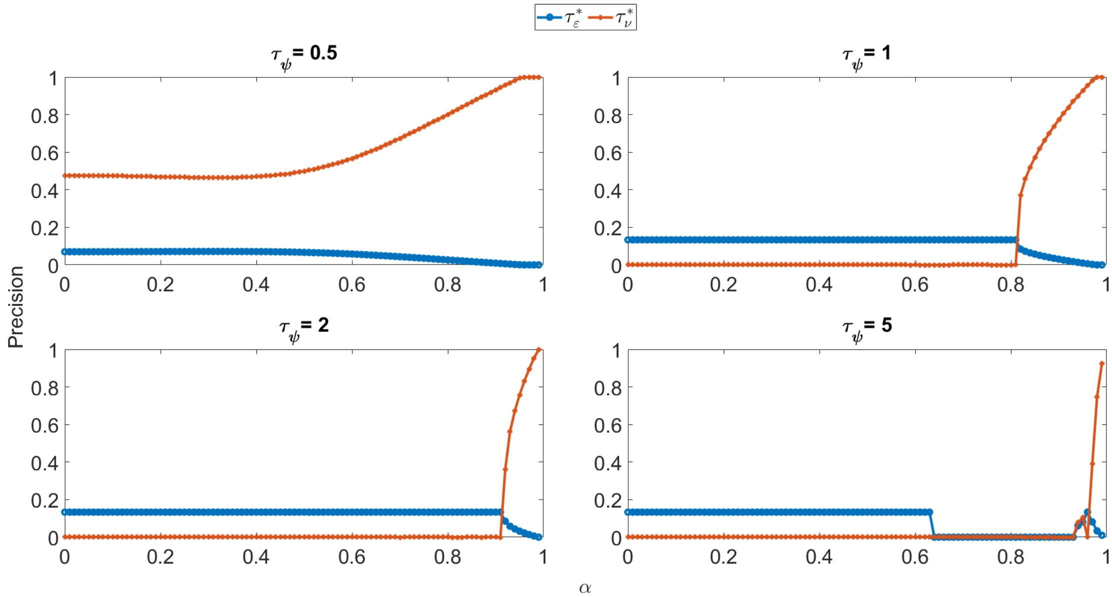

Figure 2.14. Optimal choice of τ_ν, τ_ε - $\phi = 7.5$, high τ_ψ. Each quadrant plots the optimal choices of τ_ν and τ_ε as α changes for a given level of τ_ψ. The blue line dotted with circles represents τ_ε^* and the red line dotted with pluses represents τ_ν^*. The figure uses the following fixed parameter values: $\rho = 0.95$, $\tau_w = 1$, $\phi = 7.5$, $I = 1$.

2.5.6 Summary

The optimal choice of current and revised signal precisions depends on the accuracy of private information, the degree of strategic incentives, and the budget constraint. With a small budget limit (or relatively expensive current signal precision) and imprecise private information, the S.O. maximizes welfare by providing a mixture of both signals because the gains to forecast accuracy from both signals are largest when all precisions are low. As private signal precision increases, the S.O. will put more of its budget into the current signal precision until it chooses to provide no revised signal, because more accurate private information can overcome the potential coordination costs of the current signal. However, there is a level of strategic incentives at which the S.O. will choose a mixture of both signals in order to reduce the degree of coordination on the current signal and still provide some forecasting benefits with the revised signal. This threshold increases in the budget limit (decreases in the cost of current signal precision) and the precision of private information, as both dimensions improve forecast accuracy enough to overcome coordination costs as strategic incentives rise. Finally, increasing the budget limit results in a higher degree of current precision and a lower degree of revised precision due to the lower forecasting benefits of the latter.

2.6 Application

To demonstrate the applicability of my results I conduct the following exercise. Given a set of parameters and optimal choices for τ_v and τ_ε, suppose the S.O. had the ability to either raise the budget limit I or lower the cost of current signal precision ϕ. When is it optimal to raise the budget limit instead of lowering the marginal cost, and vice versa? The answer depends on the mechanisms identified previously; namely, the impact the budget has on relative forecasting values and the degree of strategic incentives.

Figure 2.15 plots two curves depicting the relationship between welfare on the horizontal axis and either ϕ or I on the vertical axes. Both curves use the optimal choice of τ_v and τ_ε. The black solid line fixes $\phi = 5$ while the red dashed line fixes $I = 1$, as in the benchmark case, which is where the two lines necessarily cross. The figure shows that either increasing the budget limit or lowering the cost of current signal precision can increase welfare. Which choice has a greater impact on welfare depends on the specific trade-off between increasing the budget or lowering costs.

For demonstration purposes, assume the trade-off is proportional: the S.O. has the capability to either increase the budget limit by 20% or decrease ϕ by 20%. Starting from the benchmark case, that means increase I from 1 to 1.2 or lower ϕ from 5 to 4. When $I = 1.2$, welfare ≈ -1.0364, while when $\phi = 4$ welfare ≈ -1.0284, so the S.O. would prefer to lower the marginal cost rather than increase the budget limit. This is apparent from the slightly flatter slope of the red dashed line at $\phi = 5$, indicating a larger change in welfare following the same proportional change in ϕ. This result does not hold for all parameterizations; see figure 2.16 which chooses $I = 0.5, \phi = 5$ as the intersection point. Increasing the budget from $I = 0.5$ to 0.6

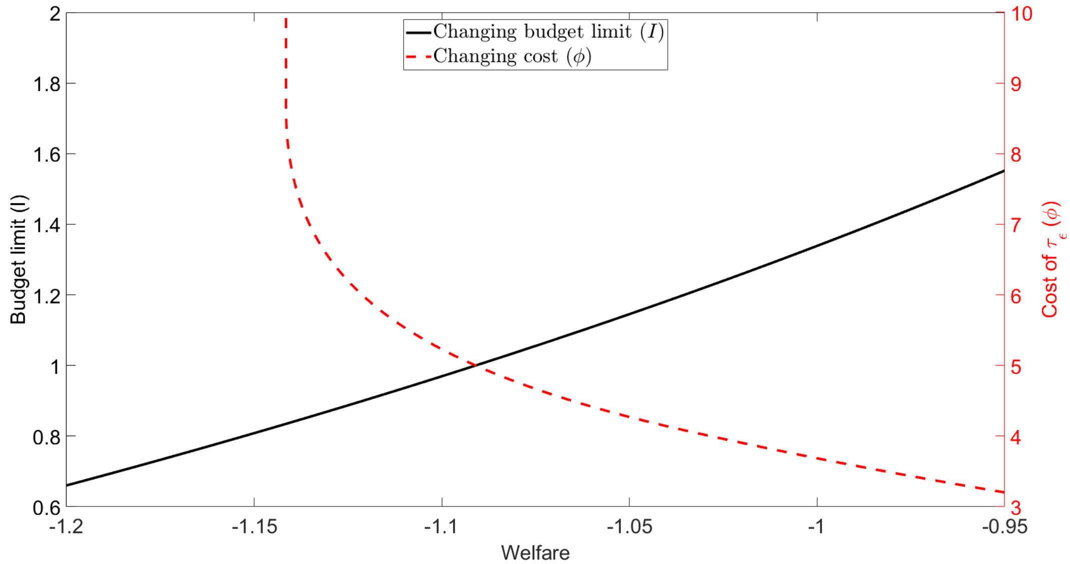

Figure 2.15. Welfare as budget limit or marginal cost changes. The black solid line depicts the level of welfare (horizontal axis) as the budget limit I changes (left-hand vertical axis) when $\phi = 5$. The red dashed line depicts the level of welfare as the marginal cost of increasing current signal precision ϕ changes (right-hand vertical axis) when $I = 1$. The maximum for both vertical axes is double the level of the benchmark case. The figure uses the following fixed parameter values: $\rho = 0.95$, $\tau_w = 1$, $\tau_\psi = 0.3$.

raises welfare by more than reducing ϕ from 5 to 4, as can be seen by the relatively steep slope of the red dashed line around the intersection.

This exercise shows there is no dominant option for the S.O. looking to raise welfare by improving its budget constraint. The slope of the welfare function with respect to the budget limit or marginal cost of precision depends on the current optimal choice for τ_ν and τ_ε, which the previous section showed depend on the degree of private precision and strategic incentives. In general, the more influential a facet of the budget is on τ_ν^* and τ_ε^*, the bigger gains to welfare from relaxing that facet. This was the example in figure 2.16 which used the same parameter values as the benchmark case except the budget limit was halved. That created the opportunity for greater welfare improvements from raising the budget instead of lowering the marginal cost of signal precision.

To maximize welfare the S.O. must therefore have accurate knowledge of many factors,

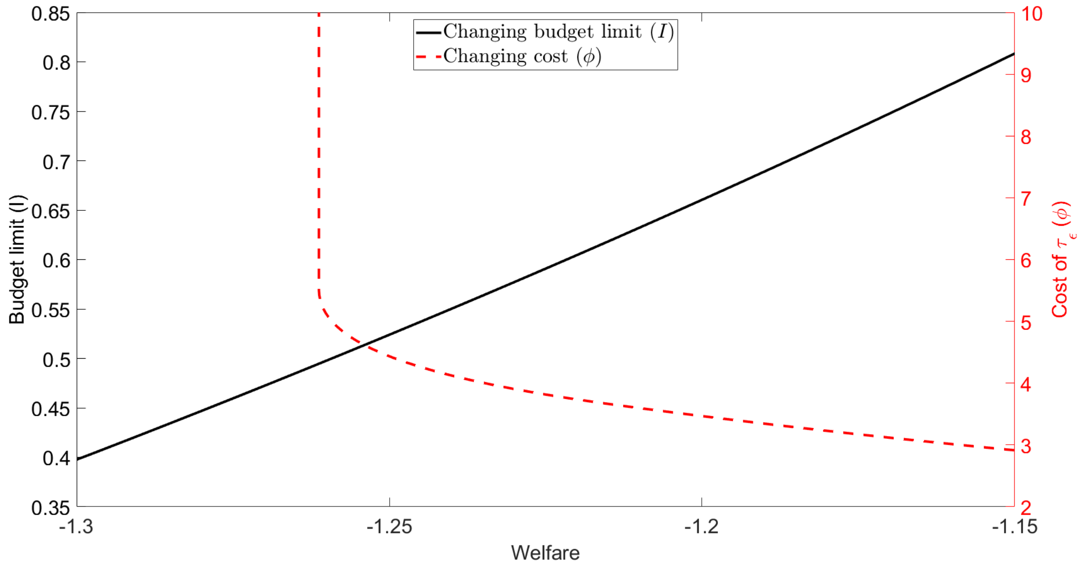

Figure 2.16. Welfare as budget limit or marginal cost changes. The black solid line depicts the level of welfare (horizontal axis) as the budget limit I changes (left-hand vertical axis) when $\phi = 5$. The red dashed line depicts the level of welfare as the marginal cost of increasing current signal precision ϕ changes (right-hand vertical axis) when $I = 0.5$. The maximum for both vertical axes is double the level of the benchmark case. The figure uses the following fixed parameter values: $\rho = 0.95$, $\tau_w = 1$, $\tau_\psi = 0.3$.

some of which may be unobservable like the properties of the fundamental (ρ and τ_w) and the precision of private information (τ_ψ). This exercise additionally requires the S.O. to precisely understand how its resources translate into signal precisions and the potential trade-offs between changes in their resource constraint. I have made simplifying assumptions to sidestep these difficulties, but my results emphasize the potential gains to optimal decision making for the S.O. which attempts to identify these parameters and relationships.

Chapter 2 is currently being prepared for submission for publication; Tyler L. Paul. The dissertation author was the principal author on this paper.

Chapter 3

Measuring Information Frictions: The Interaction of Revised Data and Higher-Order Dynamics

The main method for estimating information frictions from surveys of forecasters makes two key assumptions: the variable to be forecast follows a simple dynamic process, and forecasters observe only current data signals. I relax the first assumption by introducing additional past dependence to the fundamental. Previous authors have claimed that this requires modifying the estimation method to account for higher-order dynamics, but I prove analytically and by simulation that the existing method is still valid. The method works by relating information flows to forecast errors, thus the specific time series properties of the variable are irrelevant for estimating information frictions. When I relax the second assumption and introduce a revised data signal of the past, I find it interacts with higher-order dynamics to create additional bias in the estimation of information frictions. I propose a correction to the regression which eliminates the bias by capturing the new information flow, and I show how the correction does not depend on the specific time series properties of the data. I use simulation to study the nature of this new bias, finding it to be quantitatively small relative to the bias caused solely by revised data.

3.1 Introduction

Measuring information frictions from surveys of expectations is a simple and powerful way to recover the deep parameters necessary for modern business cycle models. The leading estimation method proposed by Coibion and Gorodnichenko (2015) (henceforth CG) exploits the predictability of average forecast errors by average forecast updates, which are changes across time of the average forecast for a given time horizon. Importantly, the CG method assumes individual forecasters only have access to data signals informative of the present time period. I show in the first chapter that the CG method suffers from omitted variable bias when forecasters also receive new signals informative of the past, i.e. revised data signals. The cause of the bias comes from the additional learning a revised signal offers which breaks the symmetry between forecast errors and forecast updates, and thus disrupts the method's ability to recover the degree of information frictions.

I proved the existence of the bias assuming the underlying economic variable to be forecast followed a simple dynamic process: an autoregression with one lag, or AR(1), meaning the current value only depends on its previous value plus an innovation term. Since revised data contain observations of past values, it is natural to wonder how learning of the past interacts with forecasting a variable that has higher-order dynamics, such as dependence on the past two or three periods. That is the goal of this paper: to understand how the revised data signal impacts measurement of information frictions when the underlying fundamental has additional persistence on the past. Specifically I answer two questions. First, how can one recover a measure of information frictions from surveys of expectations when the fundamental has higher-order dynamics? Second, how do higher-order dynamics change the bias in the original CG method when exposed to revised data?

Both questions have practical significance, since many aggregate data series in macroeconomics do not strictly follow an AR(1) as assumed in chapter one and the original CG paper. For example, when evaluating outliers in the Nelson and Plosser data set from 1982, Charles and Darné (2012) find only two of the fourteen aggregate variables are best modeled as an AR(1). Kamber, Morley and Wong (2018) demonstrate how assuming an AR(1) for the output gap can produce mistimed and unrealistically small cyclical behavior. When I separately model each data series used in chapter one I find only two out of fourteen variables are best-represented by an AR(1), with the rest exhibiting meaningful higher-order dynamics.

I address how to measure information frictions in the presence of higher-order dynamics by starting from an extension proposed by CG in their original paper. When the fundamental follows an autoregression with p lags, or AR(p), CG derive a new regression specification for measuring information frictions which involves regressing average forecast errors on the current forecast update plus a sequence of forecast updates regarding past observations of the fundamental. The additional persistence of the fundamental, they argue, requires additional forecast update terms of lagged values in order to preserve the predictability of forecast errors and thus recover information frictions from the data.

When CG test the importance of higher-order dynamics for measuring information frictions, they find the regression model which assumes an AR(1) fits the data best. This is somewhat surprising given the evidence outlined above regarding the higher-order properties for many variables, but CG argue that there is likely high correlation across past forecast updates that renders their point estimates for the new terms overly imprecise. I confirm their findings using Monte Carlo simulations of the forecasting environment, but the large sample size I use suggests a new result: the new specification is incorrect not due to lack of precision but due to near perfect multicollinearity of the regressors. Conceptually, all forecast updates are driven by the exact same learning process: the difference between the observed signal and the individual forecast of that signal. I show analytically that each forecast update term is related via the ratio of Kalman gains for the associated forecast horizon. The new specification is thus unable to reliably recover the degree of information frictions, regardless the size of the sample.

This finding adheres to the logic outlined in chapter one. Forecast updates reflect the new information learned in the current period. When that information is contaminated by noise, individuals will rationally update their previous forecasts only partially, so their new forecast is a weighted average of their prior belief and the new information. The gap between their new and old forecasts is thus entirely determined by what they learned in the current period. This applies to all forecast updates, including those regarding the fundamental in previous periods. Therefore each regressor in the new CG specification has the exact same stochastic properties, so the regression suffers from perfect multicollinearity.

It is straightforward to show analytically that the time series properties of the fundamental do not affect the correct regression specification to use when measuring information frictions. I highlight the error in CG's algebra which results from their use of the generalized inverse of a non-invertible matrix. In the end, I find the original CG method is still the correct way to measure information frictions when individuals observe one signal and the fundamental has higher-order dynamics.

With the appropriate benchmark established, I then compute the analytical results from measuring information frictions of a higher-order process when individuals use a revised signal in addition to a current signal. I find a new bias expression which consists of two sets of terms: the bias from the AR(1) case, and the sum of covariances between the forecast update (i.e. the regressor) and the forecast error made last period of the state multiple periods ago. The second set of terms is called the lagged backcast error and is present due to the interaction of revised data and the additional lags in the fundamental. Forecasting an AR(p) requires individuals to make expectations of the fundamental value in each of the previous p periods. A revised signal provides information about last period which is useful for forecasting all horizons extending backwards. The new bias term reflects individuals using information about the last period to update their forecasts of two, three, or more periods ago. This update is not captured in the current period forecast update, and thus shows up in the residual as omitted variable bias.

The nature of the fundamental process therefore does alter the structural constituents of the bias. At the same time, the nature of the fundamental process does not change the correct method for recovering information frictions. Both of these findings together imply that the correct specification with revised data is the same as in the AR(1) case, which is the same as originally suggested in chapter one: regress average forecast errors on the average forecast update plus the revised signal innovation. The latter term is new and helps capture all of the learning from the revised signal so that the coefficient on the original regressor (i.e. forecast update) can cleanly recover the degree of information frictions. This implies that the exercise conducted in chapter one was theoretically justified even though it assumed an AR(1) for all variables, since higher-order dynamics on their own do not introduce any additional bias. At the same time the original CG paper is theoretically justified given their informational framework consisting of only current period signals, despite their new specification suggesting otherwise.

However, there is additional bias due to higher-order dynamics of the process when estimating the original CG method without my correction if individuals observe revised signals. To assess the relevance of this new bias I conduct numerical simulations of the forecasting environment, varying the time series properties of the fundamental in an effort to find any patterns. I make two main findings. First, the size of the new bias is quantitatively small relative to the bias caused by revised data. The direction and magnitude of the bias depend on the persistence and propagation of the fundamental. For instance, the bias caused by higher-order dynamics is larger when the process has highly cyclical dynamics versus a similar process (in terms of unconditional variance) with smooth dynamics. Highly cyclical processes also cause the higher-order dynamics bias to be negative. That said the overall size is always small, due to a combination of heavy discounting and canceling out by successive covariance terms. Second, the bias caused by revised data decreases in the lag order of process if it has cyclical behavior. This suggests an intriguing connection between the total bias and higher-order dynamics that works by affecting the forecast update term itself. Overall, my findings suggest researchers do not need to be worried about higher-order dynamics introducing sizable bias into estimates of information frictions, regardless the specific nature of the process.

This paper is a natural follow-up to the first chapter, which sits in the literature that uses surveys of expectations to measure information frictions. This line of research gained momentum following two methodological contributions from Coibion and Gorodnichenko (2012, 2015), the former linking fundamental shocks to survey reactions while the latter connected forecast errors and forecast updates. Several authors have recently measured the expectations formation process directly through new surveys on businesses, such as McClure, Coibion and Gorodnichenko (2022), Candia et al. (2023), and Ropele, Gorodnichenko and Coibion (2024). Many of the variables targeted in these studies exhibit higher-order dynamics, making it vital to understand their interaction with all features of the information framework.

A recent trend in the literature has been to test the limits of the original CG method. A notable example is Bordalo et al. (2020) who use the CG test to identify irrational behavior by forecasters in the form of overreacting to new information. Other authors have similarly used the CG framework to argue that households do not use Bayesian updating for their beliefs, such as D'Acunto et al. (2021) who study consumer grocery purchases. The bias I identify due to higher-order dynamics could potentially reverse these results, but in practice it is too small to qualitatively make a difference. Gemmi and Valchev (2022) demonstrate that forecaster surveys exhibit bias due to strategic considerations, which results in an omitted variable bias in the CG method. While their approach is similar, the omitted variable bias I identify is due to the time series properties of the fundamental, not an assumption about individual incentives.

Finally, this paper connects to the real-time forecasting literature which can trace its modern roots to the development of the real-time data set by Croushore and Stark (2001). Many authors have used these data to create ARIMA forecasting models which use the full set of data vintages. Some prominent examples include Jacobs and van Norden (2011), Kishor and Koenig (2009), Galvão (2017), and Bańbura et al. (2013) in the handbook of forecasting. These papers demonstrate how revised data signals can improve forecasting models, but also that forecasters may not fully appreciate that fact (Clements, 2012). This paper brings the insights from the real-time forecasting literature to questions of measuring information frictions with a unique focus on the role of higher-order dynamics.

The paper proceeds as follows. First I provide motivation by estimating the time series properties of key macroeconomic variables originally studied in CG. In section 3.3 I analyze the CG specification for measuring information frictions both with and without revised data. I demonstrate why the existing method and correction term are still valid despite higher-order dynamics, and characterize the bias caused by the interaction between revised data and higher-order dynamics. Section 3.4 contains simulations which quantify the economic importance of this new bias and study its dependence on the time series properties of the fundamental. Section 3.5 concludes.

3.2 Dynamic properties of macroeconomic data

I begin by modeling the time series properties of key macroeconomic variables used in the original CG paper as well as in chapter one. Both papers assumed each variable followed an AR(1), which is a commonly-used starting point for authors disentangling trends from cyclical components (Kamber, Morley and Wong, 2018). Since this paper is concerned with the interaction between revised data and higher-order dynamics, my modeling objective is to characterize the autocorrelation structure of each macroeconomic variable in the form of an AR(p) model with statistically-identified optimal lag lengths.[1] The goal is not out-of-sample forecasting or to propose a true representation of each variable, but to motivate the importance of my research question by demonstrating the higher-order properties of many real-world variables.

I apply the modeling methodology of Box, Jenkins and Reinsel (1994) in three steps. First, I transform variables when necessary to ensure they are stationary by taking logged differences, using the augmented Dickey-Fuller test for unit roots (Dickey and Fuller, 1979). Second, I fit four autoregressive models to the data with lag lengths of $p = 1 \dots 4$. Third, I use Bayesian Information Criteria to identify the most parsimonious model with the highest likelihood given the data. I ensure all variables come from the same data vintage to avoid the impact of data revisions.

Table 3.1 contains the results. One series, real consumption growth, has optimal lag length of 0, but I choose the next-best model according to BIC to ensure compatibility with the rest of the paper. Out of the remaining 13 variables, BIC only chooses an AR(1) for two series: real GDP and housing starts. The rest exhibit some form of higher-order dynamics and often include three or four lags in the process. This underlines the importance of understanding how higher-order dynamics interact with revised data when measuring information frictions.

[1] I exclude moving average terms from the estimation to match the original CG paper.

Table 3.1. AR Model Selection for Macroeconomic Variables

	Lag length	ρ_1	ρ_2	ρ_3	ρ_4
Nominal GDP growth	4	0.41		0.13	0.14
		(0.050)		(0.058)	(0.049)
Real GDP growth	1	0.44			
		(0.044)			
GDP price inflation	3	0.51	0.13	0.27	
		(0.051)	(0.064)	(0.067)	
Real consumption growth	4	0.03	0.04	0.11	0.24
		(0.061)	(0.059)	(0.050)	(0.049)
Real nonresidential investment growth	2	0.41	0.13		
		(0.054)	(0.076)		
Real residential investment growth	2	0.59	-0.14		
		(0.038)	(0.049)		
CPI inflation	3	0.58	-0.04	0.30	
		(0.070)	(0.098)	(0.070)	
Housing starts growth	1	0.21			
		(0.056)			
Industrial production growth	2	0.65	-0.21		
		(0.047)	(0.054)		
Unemployment rate	2	1.64	-0.68		
		(0.036)	(0.037)		
Employment growth	3	-0.49	-0.47	-0.26	
		(0.086)	(0.083)	(0.093)	
Treasury bill 3 month rate	4	1.24	-0.51	0.45	-0.22
		(0.039)	(0.054)	(0.050)	(0.047)
Treasury bond 10 year rate	3	1.11	-0.27	0.14	
		(0.051)	(0.083)	(0.060)	
AAA Corporate bond rate	4	1.22	-0.39	0.34	-0.17
		(0.087)	(0.144)	(0.146)	(0.083)

Notes: Lag length determined by BIC. Standard errors are in parentheses. All variables converted to lag differences except unemployment rate, treasury bill rate, treasury bond rate, and corporate bond rate. All models estimated on quarterly data from 1965 to 2019.

3.3 Analysis of higher-order dynamics

I study an economy in which a continuum of individual forecasters use imprecise data signals to forecast a fundamental economic variable. The set-up is similar to the environment of chapter one but with a focus on higher-order dynamics.

3.3.1 Model set-up

The fundamental x_t follows an AR(p) which I represent in state-space form:

$$\text{State:} \quad \boldsymbol{z}_t \equiv \begin{bmatrix} x_t \\ x_{t-1} \\ \vdots \\ x_{t-p} \end{bmatrix} = \begin{bmatrix} \rho_1 & \rho_2 & \dots & \rho_p \\ 1 & 0 & \dots & 0 \\ 0 & \ddots & \ddots & \vdots \\ 0 & \dots & 1 & 0 \end{bmatrix} \boldsymbol{z}_{t-1} + \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} w_t$$

$$= \boldsymbol{\Phi} \boldsymbol{z}_{t-1} + \boldsymbol{S} w_t$$

where $w_t \sim N(0, \tau_w^{-1})$ is an i.i.d. shock which drives the fundamental, and $\boldsymbol{\Sigma}_w = \mathbb{E}[(Sw_t)(S'w_t)]$. Individual forecasters receive two noisy private signals, one of x_t and one of x_{t-1}:

$$\text{Measurement:} \quad \boldsymbol{y}_{it} \equiv \begin{bmatrix} y_{it} \\ r_{it} \end{bmatrix} = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \end{bmatrix} \boldsymbol{z}_t + \begin{bmatrix} \psi_{it} \\ \nu_{it} \end{bmatrix}$$

$$= \boldsymbol{H} z_t + \boldsymbol{u}_{it}$$

where $\psi_{it} \sim N(0, \tau_\psi^{-1})$ and $\nu_{it} \sim N(0, \tau_\nu^{-1})$ are i.i.d. idiosyncratic noise shocks unique to each individual, and $\boldsymbol{\Sigma}_u = \mathbb{E}[\boldsymbol{u}_{it}\boldsymbol{u}'_{it}]$.[2] All shocks are uncorrelated with each other and across time. The

[2]In chapter one, the shock ν_{it} was common across individuals to represent a data revision published by the statistical office. Changing this to an idiosyncratic shock sharpens the results regarding higher-order dynamics.

matrix of forecast errors **P** can be determined by solving the Ricatti equation:

$$\boldsymbol{P} = \boldsymbol{\Phi}(\boldsymbol{P} - \boldsymbol{P}\boldsymbol{H}'(\boldsymbol{H}\boldsymbol{P}\boldsymbol{H}' + \boldsymbol{\Sigma_w})^{-1}\boldsymbol{H}\boldsymbol{P})\boldsymbol{\Phi}' + \boldsymbol{\Sigma}_u$$

I denote $\boldsymbol{K} = \boldsymbol{P}\boldsymbol{H}'(\boldsymbol{H}\boldsymbol{P}\boldsymbol{H}' + \boldsymbol{\Sigma_w})^{-1}$ as the Kalman gain, which represents the weight individuals place on their new information when forming forecasts. The Kalman gain measures the degree of information frictions. A large Kalman gain means individuals place more weight on new information relative to their prior beliefs, implying a small degree of information frictions. For example, when the fundamental is an AR(2) the Kalman gain is a 2×2 vector: $\boldsymbol{K} = \begin{bmatrix} K_1 & K_r \\ K_2 & K_{r_2} \end{bmatrix}$. K_1 and K_2 are the weights placed on the current signal y_{it} for forecasting x_t and x_{t-1}, respectively, while K_r and K_{r_2} are the weights placed on the revised signal r_{it} for forecasting x_t and x_{t-1}, respectively.

The optimal h-period-ahead forecast of the state-vector z_{t+h} is given by:

$$\mathbb{E}_{it}[\boldsymbol{z}_{t+h}] = \mathbb{E}_{it-1}[\boldsymbol{z}_{t+h}] + \boldsymbol{K}\boldsymbol{\Phi}^h\boldsymbol{H}(\boldsymbol{y}_{it} - \mathbb{E}_{it-1}[\boldsymbol{y}_{it}]) \tag{3.1}$$

where $\mathbb{E}_{it-s}[x_{t+h}]$ is the mean-square error minimizing forecast of x_{t+h} conditional on information observed in period $t-s$ by individual i. Averaging across all individuals produces $\mathbb{E}_{t-s}[x_{t+h}]$, where I drop the i subscript for the expectations operator. I choose $h = 0$ without loss of generality in the following sections for ease of exposition.[3]

3.3.2 New CG specification for measuring information frictions

The original CG specification assumed the fundamental followed an AR(1) and forecasters only observe one current signal, which implies $\tau_v = 0$ and $\boldsymbol{H} = \begin{bmatrix} 1 & 0 & \dots & 0 \end{bmatrix}$. In this case,

[3] Solving for $h > 0$ adds additional terms in the future which are necessarily uncorrelated with information in the present, so these terms do not affect measures of information frictions or the bias.

CG suggested estimating the Kalman gain using the following regression model:

$$x_t - \mathbb{E}_t[x_t] = \beta\left(\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t]\right) + \varepsilon_t \tag{3.2}$$

Call equation (3.2) the univariate model. This model regresses average forecast errors on the average forecast update, generating estimates of β which map directly into the Kalman gain: $K_1 = \frac{1}{1+\hat{\beta}}$. When the fundamental follows an AR(p), CG suggest a new regression model:

$$\begin{aligned} x_t - \mathbb{E}_t[x_t] = \beta_{11}\left(\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t]\right) + \beta_{12}\left(\mathbb{E}_t[x_{t-1}] - \mathbb{E}_{t-1}[x_{t-1}]\right) \\ + \cdots + \beta_{1p}\left(\mathbb{E}_t[x_{t-(p-1)}] - \mathbb{E}_{t-1}[x_{t-(p-1)}]\right) + \varepsilon_t \end{aligned} \tag{3.3}$$

Call equation (3.3) the multivariate model. This model directly nests the univariate one, the difference being $(p-1)$ new forecast update terms regarding the fundamental in previous periods: $x_{t-1}, x_{t-2}, \ldots, x_{t-(p-1)}$. The authors show the estimated coefficients from this model also map directly into the Kalman gains:[4]

$$\begin{aligned} \beta_{11} &= \frac{K_1}{K_1^2 + K_2^2 + \cdots + K_p^2} - 1 \\ \beta_{12} &= \frac{K_2}{K_1^2 + K_2^2 + \cdots + K_p^2} \\ \ldots\ \beta_{1p} &= \frac{K_p}{K_1^2 + K_2^2 + \cdots + K_p^2} \end{aligned}$$

At face value both of these properties are appealing and intuitive. However, the new forecast updating terms are only included due to an incorrect algebra step. To show this I follow

[4]The online appendix of CG incorrectly writes the denominators as $(K_1 + \cdots + K_p)^2$.

the CG procedure outlined in their appendix starting from equation (3.1):[5]

$$(\boldsymbol{I}-\boldsymbol{KH})(\mathbb{E}_t[\boldsymbol{z}_t]-\mathbb{E}_{t-1}[\boldsymbol{z}_t]) = \boldsymbol{KH}(\boldsymbol{z}_t - E_t[\boldsymbol{z}_t]) \tag{3.4}$$

This equation relates forecast errors to forecast updates, similar to regression specifications (3.2) and (3.3). To isolate forecast errors as the dependent variable requires taking the inverse of $\boldsymbol{KH}$, but this square matrix is not invertible. For an AR(2), this looks like

$$\boldsymbol{KH} = \begin{bmatrix} K_1 \\ K_2 \end{bmatrix} \begin{bmatrix} 1 & 0 \end{bmatrix} = \begin{bmatrix} K_1 & 0 \\ K_2 & 0 \end{bmatrix}$$

Instead CG take the generalized inverse of $\boldsymbol{KH}$ which they claim produces the following:

$$\boldsymbol{z}_t - \mathbb{E}_t[\boldsymbol{z}_t] = \{\boldsymbol{KH}\}^{+}(\boldsymbol{I}-\boldsymbol{KH})(\mathbb{E}_t[\boldsymbol{z}_t]-\mathbb{E}_{t-1}[\boldsymbol{z}_t])$$

$$\begin{bmatrix} x_t - E_t[x_t] \\ x_{t-1} - E_t[x_{t-1}] \\ \vdots \\ x_{t-p} - E_t[x_{t-p}] \end{bmatrix} =$$

$$(K_1^2 + K_2^2 + \dots)^{-1} \begin{bmatrix} K_1 & \dots & K_p \\ 0 & \dots & 0 \\ \vdots & \ddots & 0 \\ 0 & \dots & 0 \end{bmatrix} \begin{bmatrix} 1-K_1 & 0 & \dots & 0 \\ -K_2 & 1 & \ddots & 0 \\ \vdots & 0 & \ddots & 0 \\ -K_p & 0 & \dots & 1 \end{bmatrix} \begin{bmatrix} \mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t] \\ \mathbb{E}_t[x_{t-1}] - \mathbb{E}_{t-1}[x_{t-1}] \\ \vdots \\ \mathbb{E}_t[x_{t-(p-1)}] - \mathbb{E}_{t-1}[x_{t-(p-1)}] \end{bmatrix}$$

The top row of this system of equations is equation (3.3). However, the entire system cannot be correct since every other row implies there is no forecast error for past variables; for example, the value of the second row is $x_{t-1} - \mathbb{E}_t[x_{t-1}] = 0$. In reality there should be some error of the past since it is never perfectly revealed. I demonstrate this by deriving the true expression for

[5] First plug in $\boldsymbol{y}_{it} = \boldsymbol{Hz}_t$, and then notice that $\mathbb{E}_{it-1}[\boldsymbol{u}_{it}] = 0$. Next average both sides across all individuals so $\int \boldsymbol{u}_{it} di = 0$. Finally subtract $\boldsymbol{KH}E_t[\boldsymbol{z}_t]$ from both sides and rearrange.

$x_{t-1} - \mathbb{E}_t[x_{t-1}]$ when $p = 2$ in appendix A.2.2.[6]

The problem with this procedure comes from taking the generalized inverse of $\boldsymbol{KH}$ which in general has different properties than the inverse of a non-singular matrix. I avoid this issue by expanding out equation (3.4) without taking the generalized inverse:

$$\begin{bmatrix} x_t - \mathbb{E}_t[x_t] \\ x_{t-1} - \mathbb{E}_t[x_{t-1}] \\ \vdots \\ x_{t-p} - \mathbb{E}_t[x_{t-p}] \end{bmatrix} = \begin{bmatrix} \frac{1-K_1}{K_1}(\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t]) \\ K_2^{-1}(\mathbb{E}_t[x_{t-1}] - \mathbb{E}_{t-1}[x_{t-1}]) - (\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t]) \\ \vdots \\ K_p^{-1}(\mathbb{E}_t[x_{t-(p-1)}] - \mathbb{E}_{t-1}[x_{t-(p-1)}]) - (\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t]) \end{bmatrix} \tag{3.5}$$

Equation (3.5) reveals the correct relationship between forecast errors and forecast updates for an AR(p). First note that past forecast errors are non-zero and depend on current plus lagged forecast updates.[7] More importantly, the first row matches the univariate model exactly, both by the coefficient mapping and the number of regressors. This means that even when the fundamental has higher-order dynamics the univariate model is the correct method to measure information frictions. When the fundamental has any number of lags p, current forecast errors are only predictable by the current forecast update.

The dynamic properties of the fundamental do not change the ability of the univariate model to measure information frictions. Intuitively this is because the univariate model works by evaluating the usage of new information in forecasts. The dependent variable contains the portion of the fundamental which the average forecast misses. This is regressed on the forecast update, which contains how individuals incorporate new information into their forecasts. Both variables are driven by the same stochastic process, the surprise in the observed signal, so the

[6] The expression in the appendix includes a revised data signal. Simply set $K_r = K_{r_2} = 0$ to find the expression with only one current signal.

[7] The goal of this exercise was to demonstrate the structural relationship between forecast errors and updates, so the presence of time t information on the right-hand side is not an issue. It does, however, mean the relationship is non-causal and cannot be implemented in real-time. Later I will derive a regression specification that can in theory be used in real-time to recover $K_2, \dots K_p$.

regression contains a tight mapping back to the Kalman gain. This is purely an informational mechanism; it reflects how individuals use data signals to update their forecasts of the state, not the underlying properties of the state itself. If the state is composed of many lags as in the AR(p) case, individuals will update their forecasts of each lag according to the same learning from their signals. The different Kalman gain parameters $(K_1, K_2, \ldots)$ capture how that same signal observation affects each horizon's forecast $(x_t, x_{t-1}, \ldots)$.

To measure these gain parameters I propose the following method.

Proposition 4 (Measuring higher-order Kalman gains) *When the fundamental follows an AR(p) and forecasters observe one data signal of the present, the Kalman gain parameter K_s for $s \in (2, 3, \ldots, p)$ can be estimated from the following regression:*

$$\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t] = \beta_s(\mathbb{E}_t[x_{t-(s-1)}] - \mathbb{E}_{t-1}[x_{t-(s-1)}]) + \varepsilon_t \tag{3.6}$$

where $K_s = \frac{K_1}{\hat{\beta}_s}$, and K_1 is estimated from the univariate model (3.2).

Call equation (3.6) the univariate+ model. The proof for this proposition comes from expanding out equation (3.4) and noticing that the forecast update for each forecast horizon $s \in (0, 1, \ldots)$ has a similar structure:

$$\mathbb{E}_t[x_{t-(s-1)}] - \mathbb{E}_{t-1}[x_{t-(s-1)}] = K_s(x_t - \mathbb{E}_{t-1}[x_t]) \tag{3.7}$$

This precisely shows that all average forecast updates are driven by the same learning process, $x_t - \mathbb{E}_{t-1}[x_t]$. To get the equation in the proposition, first set $s = 1$ in equation (3.7) to find $\frac{1}{K_1}\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t] = (x_t - \mathbb{E}_{t-1}[x_t])$. Plug this into the right-hand side of equation (3.7) to find equation (3.6) from the proposition. To run this regression on survey data, individual forecasters must report their current forecasts of the variable s-periods in the past, which in practice is not often reported.

However, I can numerically simulate the model to confirm the accuracy of proposition 4 as well as the claim that the univariate model will always recover K_1 while the multivariate model will not always recover K_s for $s \in (1,2,\dots)$. I create 500 draws of the fundamental, noisy signals for each individual, and their optimal forecasts across time. I compute the average forecast error and forecast updates for x_t and x_{t-1} for each draw of the economy, and then I estimate the two univariate regressions (3.2), (3.6), and the multivariate regression (3.3). I compare the average implied Kalman gain from each model with the true value as dictated by the steady-state Kalman filter.

Table 3.2 shows results from simulating an AR(2) using different autoregression parameter values. In general, the two univariate models (original CG specification (3.2) for K_1 and the model from proposition 2 (3.6)) almost always perfectly recover the true values of K_1 and K_2. This supports the claim that estimating information frictions does not require additional terms due to higher-order dynamics, since the univariate models entirely capture the effect of information flows. The multivariate model performs significantly worse for many parameterizations of the fundamental. For instance, when $(\rho_1, \rho_2) = (1.4, -0.5)$, the multivariate estimates $K_1 = 0.97$ when the true value is 0.77, a bias of 20 percentage points. However, this is not always the case; when $(\rho_1, \rho_2) = (0.1, 0.65)$, the estimated K_1 from the multivariate model is very near to the true value. The same is not true for estimates of K_2 produced by the multivariate model, which are rarely close to the true value and often wildly inaccurate.

In summary, the time series properties of the fundamental do not change the method for measuring information frictions from forecasting data. The univariate model still applies for any AR(p) because it directly measures how new information is used. When individuals only observe one data signal, they update their forecasts of all horizons using the information in that signal, so including additional forecast update terms is not necessary. Doing so can produce inaccurate estimates of information frictions. Next I demonstrate how the time series properties of the fundamental introduce bias to the univariate model when individuals observe a revised signal alongside their current signal.

Table 3.2. Comparing univariate and multivariate models of information frictions - AR(2)

		Estimate of K_1			Estimate of K_2		
ρ_1	ρ_2	**True**	**Uni**	**Multi**	**True**	**Uni**	**Multi**
1.90	-0.90	0.82	0.82	1.08	0.24	0.24	-0.65
1.40	-0.90	0.78	0.78	0.92	0.20	0.20	-0.36
1.00	-0.90	0.75	0.75	0.81	0.15	0.15	-0.14
0.50	-0.90	0.73	0.73	0.72	0.08	0.08	0.15
1.40	-0.50	0.77	0.77	0.97	0.22	0.22	-0.47
0.90	-0.20	0.72	0.72	0.81	0.17	0.17	-0.22
0.90	0.09	0.72	0.72	0.83	0.18	0.19	-0.24
0.60	-0.80	0.72	0.72	0.73	0.10	0.10	0.04
0.60	-0.20	0.69	0.69	0.74	0.12	0.12	-0.14
0.60	0.35	0.71	0.71	0.76	0.14	0.14	-0.15
0.30	-0.90	0.72	0.72	0.70	0.05	0.05	0.39
0.30	-0.50	0.69	0.69	0.68	0.06	0.06	0.12
0.30	-0.20	0.68	0.68	0.68	0.06	0.06	-0.03
0.30	0.35	0.68	0.68	0.70	0.07	0.07	-0.08
0.30	0.65	0.70	0.70	0.71	0.08	0.08	-0.01
0.10	-0.90	0.72	0.72	0.69	0.02	0.02	1.39
0.10	-0.50	0.69	0.69	0.67	0.02	0.02	0.51
0.10	-0.20	0.67	0.67	0.67	0.02	0.02	0.03
0.10	0.35	0.68	0.68	0.68	0.02	0.02	-0.15
0.10	0.65	0.70	0.70	0.69	0.03	0.03	0.16
0.10	0.85	0.72	0.72	0.65	0.03	0.03	1.82

Notes: True K_s refers to the Kalman gain parameter derived from the steady-state Kalman filter. Univariate refers to model (3.2) for K_1 and model (3.6) for K_2, K_3, K_4. Multivariate refers to model (3.3). The values in each of these columns are the average implied estimate of the Kalman gain across all simulations given that row's autoregression model. Each simulated economy had 750 individuals and 750 time periods. Each AR model had 500 economies generated. I fix $\tau_w = 1, \tau_\psi = 2$.

3.3.3 Measurement bias: revised data and higher-order dynamics

I reintroduce the revised signal by letting $\tau_v > 0$ and analyze how the interaction of revised data and higher-order dynamics introduces a novel bias to measurements of information frictions. The previous section showed the method for measuring information frictions when the fundamental is an AR(p) is the same as when the fundamental is an AR(1), which means using the original CG univariate model (3.2). In that case, the structural relationship between average forecast errors and average forecast updates is:

$$x_t - \mathbb{E}_t[x_t] = \left(\frac{1-K_1}{K_1}\right)(\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t])$$

If individual forecasters also observe a revised data signal, this structural relationship changes according to proposition 5:

Proposition 5 (Structural equation with revised data) *When the fundamental is an AR(p) and forecasters observe two data signals, one of the present and one of the past, the true structural relationship between average forecast errors and average forecast updates is:*

$$\begin{aligned} x_t - \mathbb{E}_t[x_t] &= \left(\frac{\rho_1(1-K_1)-K_r}{\rho_1 K_1 + K_r}\right)(\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t]) \\ &+ \left(\frac{\rho_2 K_r}{\rho_1 K_1 + K_r}\right)(x_{t-2} - \mathbb{E}_{t-1}[x_{t-2}]) + \left(\frac{\rho_3 K_r}{\rho_1 K_1 + K_r}\right)(x_{t-3} - \mathbb{E}_{t-1}[x_{t-3}]) \\ &+ \cdots + \left(\frac{\rho_p K_r}{\rho_1 K_1 + K_r}\right)(x_{t-p} - \mathbb{E}_{t-1}[x_{t-p}]) + \left(\frac{K_r}{\rho_1 K_1 + K_r}\right) w_t \end{aligned}$$

The proof is in appendix A.2.1. There are two sets of new terms on the right-hand side relative to when forecasters only observe the current signal. The first set consists of the final term, w_t, which is the current period fundamental. This term is also present in the structural bias equation from chapter one when the fundamental follows an AR(1), so it is not due to higher-order dynamics. The second set consists of $p-1$ terms: $(x_{t-2} - \mathbb{E}_{t-1}[x_{t-2}]), (x_{t-3} - \mathbb{E}_{t-1}[x_{t-3}]), \ldots, (x_{t-p} - \mathbb{E}_{t-1}[x_{t-p}])$. Each term is the forecast error made last period ($t-1$) when forecasting the fundamental value $2,\dots,p$ periods ago, also known as lagged backcast errors. These terms are new relative to the AR(1) case and reflect the unique interaction of revised data with higher-order dynamics. Note that if either $\rho_2 = \rho_3 = \cdots = \rho_p = 0$ (i.e. an AR(1)) or $K_r = 0$ (i.e. an irrelevant revised signal), all lagged backcast errors disappear. For ease of exposition I assume $p = 2$ for the rest of this section. In that case, $\rho_3,\dots,\rho_p = 0$ and there would only be a single lagged backcast error in proposition 5: $x_{t-2} - \mathbb{E}_{t-1}[x_{t-2}]$.

The presence of these two sets of terms in proposition 5 means trying to estimate regression specification (3.2) will lead to biased estimates of β. Specifically the bias is given by the normalized covariance between these terms and the forecast update. In the AR(2) case with only a single lagged backcast error, this covariance looks like:

$$\frac{\mathbb{E}\left[(E_t[x_t] - E_{t-1}[x_{t-1}])(x_{t-2} - E_{t-1}[x_{t-2}])\right]}{\mathbb{E}[E_t[x_t] - E_{t-1}[x_t]]^2} + \frac{\mathbb{E}\left[(E_t[x_t] - E_{t-1}[x_{t-1}])w_t\right]}{\mathbb{E}[E_t[x_t] - E_{t-1}[x_t]]^2} \tag{3.8}$$

There are two forces biasing the estimates from the original CG regression. The first force represented by the first term in (3.8) comes from the interaction of higher-order dynamics with the revised signal observation. The second force comes from the novel learning granted by the revised signal that is not represented in the forecast update. The revised signal grants individuals an additional noisy measurement of the state in the previous period, which they will use to update their forecasts. However, the additional weight forecasters place on the past state is not captured by the average forecast error and thus ends up in the regression residual. When the revised signal noise is idiosyncratic it washes out in the aggregate which makes this biasing force strictly positive. If the signal noise was common across individuals, like in chapter one, this term would only be positive if the revised signal is sufficiently precise relative to the current signal.

To understand where the lagged backcast term comes from I compute the ARMA representation of the optimal forecast in appendix A.2.2:

$$\mathbb{E}_t[x_t] = \frac{K_1 x_t + (\rho_2 K_2 + K_r)x_{t-1} + \rho_2(K_{r_2}(1-K_1) + K_2 K_r)(x_{t-2})}{1+\phi_1 L + \phi_2 L^2}$$

where $\phi_1 = (\rho_2 K_2 - \rho_1(1-K_1) + K_r)$ and $\phi_2 = \rho_2(K_2 K_r - (1-K_1)(1-K_{r_2}))$. I identify the interaction of revised data and higher-order dynamics by finding terms with $\rho_2 K_r$ and $\rho_2 K_{r_2}$, just as in the coefficient on the bias term in proposition 5. This shows up in the coefficient on x_{t-2}: $\rho_2(K_{r_2}(1-K_1) + K_2 K_r)$. This term constitutes the additional weight in the forecasting equation due to higher-order dynamics interacting with the revised signal. Whereas terms in K_r alone represent pure informational effects (i.e. learning about the past), terms in $\rho_2 K_r$ reflect using the information in the revised signal to improve backcasts. The forecast thus has additional weight on x_{t_2}, beyond that from the geometric sequence in lag polynomials, due to learning about x_{t-1} and needing to forecast x_{t-2}. When I compute the forecast error, these terms are absent:

$$x_t - \mathbb{E}_t[x_t] = \frac{(1-K_1)(1-\rho_1 L - \rho_2 L^2)x_t}{1+\phi_1 L + \phi_2 L^2}$$

There is an additional moving average term in $\rho_2 x_{t-2}$ due to the AR(2) nature of the state, but this is unrelated to the new learning provided by the revised signal since K_r and K_{r_2} are absent. This discrepancy between the average forecast (and thus forecast update) and the forecast error is the reason the lagged backcast error appears in the structural relationship. The revised signal provides additional information about past values which forecasters use to update their backcasts; the forecast update ignores these backcast updates, so the relationship between errors and updates suffers from omitted variable bias.

There is an important subtlety regarding the subject of the bias. The bias described in equation (3.8) results from estimating the univariate model (3.2) when the true structural equation is given in proposition 5. However, this bias is relative to the true structural coefficient, $\bar{\beta} = \frac{\rho_1(1-K_1)-K_r}{\rho_1 K_1 + K_r}$, and not to the structural coefficient of the univariate model, $\frac{1-K_1}{K_1}$. This is because the presence of revised data changes the structural relationship between forecast errors and forecast updates. With that in mind, the impact of the bias on estimates of K_1 is described below:

$$K_1 = \frac{1}{1+\hat{\beta}} - \frac{K_r}{\rho_1} \tag{3.9}$$

where $\hat{\beta}$ is the estimated coefficient from regression (3.2) when the true structural relationship is given by proposition 5. When the bias is positive and $\hat{\beta}$ is larger this will tend to lower the estimate of K_1. I will use equation (3.9) in section 3.4 to numerically calculate the impact of the bias on estimates of K_1.

3.3.4 Bias correction for AR(p)

The two sets of biasing terms in proposition 5 can be understood as omitted variables in the univariate regression model. However, since both sets of terms are present due to additional learning from the revised signal, I propose to correct the bias by including a new regressor that captures this new flow of information. Specifically, adding the difference between the average forecast of the revised signal and its realization will isolate the knowledge gained from the new signal and completely remove the bias. This is the same correction I proposed in chapter one for when the fundamental follows an AR(1). The corrected regression specification is as follows:

$$x_t - \mathbb{E}_t[x_t] = \beta(\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t]) + \beta_r(r_{it} - \mathbb{E}_{t-1}[x_{t-1}]) + \varepsilon_t \tag{3.10}$$

The correction is valid when the fundamental follows an AR(p) for the same reason that the original CG specification is valid for an AR(p) without revised data. The regression works by relating forecast errors to how individuals incorporate new information in their forecasts. By including the revised signal innovation as an additional regressor, all of the new learning from the revised signal is now captured and won't form part of the regression residual. The amount of lags in the fundamental is irrelevant since the regression solely operates on the flow of information in order to estimate information frictions.

3.4 Simulations

To understand the relevance of the bias caused by higher-order dynamics, I simulate economies of many individuals using their signals to forecast an economic variable. Specifically, I generate the fundamental according to a given autoregressive model for T periods. I create I noisy signals of the fundamental in the current period (y_{it}) and in the previous period (r_{it}), and compute forecasts for each individual $i \in I$ given their signals. I take averages across individual forecast updates and forecast errors and estimate the original CG regression (3.2). I use the coefficient estimate $\hat{\beta}$ to compute the implied Kalman gain $\hat{K}_1$, which I compare to the true Kalman gain as determined from the steady-state Kalman filter. To improve statistical accuracy I average these estimates across many simulated economies.

3.4.1 Benchmark results

To begin I fix the following parameters: $\tau_w = 1, \tau_\psi = 0.5, \tau_\nu = 2$. These parameters imply the current signal is four times as noisy as the revised signal, which is a reasonable starting point given the enhanced precision of data revisions as highlighted in chapter 2. Table 3.3 shows the results of simulating four economies with different AR(p) models. The second row describes the average bias across simulations of estimating K_1 from regression (3.2), computed as the difference between the estimate $\hat{K}_1$ and the actual K_1. Rows three and four decompose this bias into the two forces described by proposition 5: bias from observing the revised signal (row 3) and bias from the interaction of the revised signal and higher-order dynamics (row 4). Row three is computed as the covariance of the regressor with the fundamental w_t, while row four is computed as the covariance of the regressor with the sum of lagged backcast errors $\sum_{j=2}^{p}(x_{t-j} - \mathbb{E}_{t-1}[x_{t-j}])$, where p is given by the column number. The model in the first column is an AR(1) which has no bias due to higher-order dynamics, though I still calculate the covariance with the lagged backcast error term, $x_{t-1} - \mathbb{E}_{t-1}[x_{t-1}]$, as a validation check.

Table 3.3. Bias decomposition: benchmark parameters

	AR(1)	**AR(2)**	**AR(3)**	**AR(4)**
True K_1	0.36	0.39	0.39	0.41
Total bias	-0.286	-0.248	-0.214	-0.206
Bias from revisions	-0.287	-0.242	-0.207	-0.194
Bias from dynamics	0.000	-0.002	-0.008	-0.008
ρ_1	0.60	0.60	0.60	0.60
ρ_2	-	-0.80	-0.80	-0.80
ρ_3	-	-	0.60	0.60
ρ_4	-	-	-	-0.80

Notes: Total bias is the difference between the true K_1 and the estimated $\hat{K}_1$ produced by the original CG specification (3.2). Bias from revisions is the normalized covariance between the average forecast update and the fundamental shock w_t multiplied by $\left(\frac{K_r}{\rho_1 K_1 + K_r}\right)$. Bias from dynamics is the sum of normalized covariances between the average forecast update and lagged backcast errors multiplied by $\left(\frac{\rho_s K_r}{\rho_1 K_1 + K_r}\right)$ for $s \in (2, \dots, p)$. Each simulated economy had 750 individuals and 750 time periods. Each AR model had 100 economies generated; the bias figures are averages across all economies.

For the benchmark parameters, the CG specification produces overall negatively-biased estimates of the Kalman gain, finding stronger information frictions (i.e. smaller Kalman gains) than actually exist.[8] The vast majority of this bias comes from the revised signal, which is associated with the w_t term in equation (3.8) and was the subject of chapter one. Higher-order dynamics are responsible for a very small portion of the overall bias on K_1. As an example, the bias in K_1 caused by higher-order dynamics is less than 5% of the total bias for an AR(4).

The bias to the coefficient estimate of β_1 (not shown) is larger than the bias to the Kalman gain estimate, but its impact on K_1 is mitigated by the coefficient on the lagged backcast error term from proposition 5: $\frac{\rho_p K_r}{\rho_1 K_1 + K_r}$. In particular, K_r is always much smaller than K_1; under benchmark parameters and the AR(4) model, $K_r = 0.15$ compared to $K_1 = 0.41$. Therefore

[8]Recall the condition from chapter one to determine if the coefficient estimate $\hat{\beta}$ will be positively biased: $\frac{\tau_v}{\tau_w} > \rho_1 \frac{K_r}{K_1}$. Since this is true in table 3.3, we'd expect the bias on $\hat{K}$ to be negative according to equation (3.9).

the impact of bias to the coefficient estimate is severely reduced when applied to K_1. This is not surprising, since the average forecast update already contains most of the new information learned in period t, so the influence of lagged backcast errors must be small.

Interestingly, increasing the lag order of the fundamental lowers the magnitude of the total bias, from -0.286 for an AR(1) to -0.206 for an AR(4). Since the third row is quantitatively small across all models, this must be due to a diminishing covariance between average forecast errors and the fundamental w_t, which we see across the fourth row. This means the additional learning from the revised signal is less important when estimating information frictions when the fundamental has higher-order dynamics. Intuitively, the signal of x_{t-1} provides less information about x_t when the additional lags are heavily weighted, so the resulting bias should be lower. The benchmark parameters use large weights for each lag which is why we see decreasing bias as we add lags to the model.[9]

Another way to consider how the relative weights for fundamental lags impact the bias is by assessing the roots of the process's characteristic equation. The autoregressive parameters chosen for table 3.3 dictate covariance-stationary processes that have at least one pair of complex conjugate roots for each model.[10] As a result these processes exhibit sinusoidal impulse response functions, varying above and below 0 with dampening amplitude over time (Hamilton, 1994). To assess whether the shape of the process affects the bias caused by higher-order dynamics or not I repeat the simulation exercise but for a set of parameters that guarantee a covariance-stationary process with at least one real root at all lag lengths. Table 3.4 shows the results.[11] When comparing across tables, the models with real roots display smooth exponential decay in response to a shock, while the complex roots of table 3.3 generate an oscillating response over time.

Under the new parameterization that ensures real roots, higher-order dynamics make

[9] Larger weights also increase the unconditional variance of the process, which I examine in the following section.

[10] The AR(4) model has two pairs of complex conjugate roots, while the AR(3) has one real root and two complex roots.

[11] There are two real roots for the AR(2), one real root for the AR(3), and two real roots for the AR(4).

Table 3.4. Bias decomposition: smooth impulse response

	AR(1)	**AR(2)**	**AR(3)**	**AR(4)**
True K_1	0.36	0.36	0.36	0.36
Total bias	-0.286	-0.293	-0.295	-0.296
Bias from revisions	-0.287	-0.294	-0.296	-0.298
Bias from dynamics	0.000	0.004	0.005	0.007
ρ_1	0.60	0.60	0.60	0.60
ρ_2	-	0.10	0.10	0.10
ρ_3	-	-	0.05	0.05
ρ_4	-	-	-	0.10

Notes: Total bias is the difference between the true K_1 and the estimated $\hat{K}_1$ produced by the original CG specification (3.2). Bias from revisions is the normalized covariance between the average forecast update and the fundamental shock w_t multiplied by $\left(\frac{K_r}{\rho_1 K_1 + K_r}\right)$. Bias from dynamics is the sum of normalized covariances between the average forecast update and lagged backcast errors multiplied by $\left(\frac{\rho_s K_r}{\rho_1 K_1 + K_r}\right)$ for $s \in (2, \ldots, p)$. Each simulated economy had 750 individuals and 750 time periods. Each AR model had 100 economies generated; the bias figures are averages across all economies.

an even smaller contribution to the total bias despite a larger total bias. Total bias no longer decreases as lag length rises, the opposite result from table 3.3 but as expected given the relatively more valuable revised signal for a process with smooth exponential decay. Finally, the impact of higher-order dynamics on the bias has turned positive, which reflects the positive lag weights. Overall, higher-order dynamics can materially affect the bias only when the process has strongly cyclical behavior. Variables with many lags exhibit greater total bias, but this comes from the revised signal itself, and not its interaction with higher-order dynamics.

3.4.2 Unconditional variance of the process

The previous section showed higher-order dynamics make a quantitatively small contribution to the overall bias, though the precise magnitude depends on the cyclical properties of the fundamental. To address this more specifically I conduct two exercises. First, I compare the bias contribution across all new parameterizations with increasing unconditional variance as determined by the autoregression parameters. Higher unconditional variance reduces the value of the revised signal relative to the current signal, since the latter has better predictive power. In turn this should lower the bias caused by the revised signal. To see this result in a simpler environment, table 3.5 simulates an AR(1) with increasing values for ρ_1. Raising ρ_1 causes the unconditional variance to rise: $\frac{\tau_w^{-1}}{1-\rho_1^2}$. As the degree of persistence increases, the total bias (and thus bias caused by the revised signal) decreases in size. The rise in unconditional variance is matched by a rise in K_1, which is partly responsible for lowering the total bias as individuals shift weight from the revised signal to the current signal.

Table 3.5. Bias as AR parameters vary, AR(1)

ρ_1	**True K_1**	**Total Bias**
0.2	0.34	-0.37
0.3	0.34	-0.35
0.4	0.34	-0.33
0.5	0.35	-0.31
0.6	0.36	-0.29
0.7	0.36	-0.26
0.8	0.37	-0.24
0.9	0.38	-0.23

Notes: Total bias is the difference between the true K_1 and the estimated $\hat{K}_1$ produced by the original CG specification (3.2). Each row had 100 economies simulated, with 750 individuals and 750 time periods. The bias figures are averages across all simulations.

For the second exercise I vary the autoregression parameters in a way that preserves the unconditional variance of the process, to asses if the bias contribution changes more in response to one lag parameter versus any other. I simulate an AR(2) in table 3.6 by picking ρ_1 and ρ_2 to target three unconditional variances. The first block has the variance of the process from table 3.3, the second block has the variance of the process from table 3.4, and the variance of the third block is arbitrarily chosen as roughly double the variance of the second block.

As the unconditional variance rises the total bias falls, evident by comparing the bias in the bottom row of each block. At the same time, within each block the total bias falls as ρ_1 rises.

Table 3.6. Bias as AR parameters vary, AR(2)

Variance	ρ_1	ρ_2	True K_1	Bias Total	Bias Revisions	Bias Dynamics
1.8	0.2	-0.67	0.36	-0.311	-0.307	0.005
1.8	0.5	0.17	0.36	-0.311	-0.311	0.006
1.8	0.6	-0.60	0.37	-0.255	-0.252	-0.006
1.8	0.6	0.10	0.36	-0.294	-0.295	0.004
1.8	0.8	-0.49	0.38	-0.226	-0.223	-0.010
1.8	0.8	-0.40	0.38	-0.222	-0.220	-0.011
3.1	0.2	0.65	0.36	-0.432	-0.428	0.021
3.1	0.6	-0.80	0.39	-0.248	-0.243	-0.002
3.1	0.8	0.04	0.37	-0.249	-0.250	0.002
3.1	0.9	-0.75	0.40	-0.200	-0.195	-0.012
3.1	1.1	-0.63	0.41	-0.173	-0.169	-0.017
3.1	1.1	-0.48	0.41	-0.174	-0.171	-0.016
8.0	0.4	-0.93	0.39	-0.274	-0.267	0.009
8.0	1.0	-0.06	0.39	-0.208	-0.209	-0.003
8.0	1.2	-0.27	0.41	-0.175	-0.173	-0.012
8.0	1.3	-0.88	0.44	-0.154	-0.147	-0.020
8.0	1.3	-0.41	0.43	-0.157	-0.154	-0.017
8.0	1.5	-0.77	0.45	-0.136	-0.129	-0.024

Notes: Variance is the unconditional variance of the process given AR parameters in columns 1 and 2. Total bias is the difference between the true K_1 and the estimated $\hat{K}_1$ produced by the original CG specification (3.2). Bias from revisions is the normalized covariance between the average forecast update and the fundamental shock w_t multiplied by $\left(\frac{K_r}{\rho_1 K_1 + K_r}\right)$. Bias from dynamics is the sum of normalized covariances between the average forecast update and lagged backcast errors multiplied by $\left(\frac{\rho_s K_r}{\rho_1 K_1 + K_r}\right)$ for $s \in (2, \ldots, p)$. Each row had 300 economies simulated, with 1000 individuals and 1000 time periods. The bias figures are averages across all simulations.

Both of these results match the AR(1) case: higher unconditional variance and more first-order persistence lower the total bias. On the other hand, the bias caused by higher-order dynamics appears to rise as the unconditional variance increases. When the variance is 1.8, the average absolute value of the bias is 0.007, or 2.6% of the total, while this doubles to 0.014, or nearly 8% of the total when the variance rises to 8.0. This opposing force thus slightly elevates the total bias for higher variance processes, but the overall magnitude is still quite small compared to the contribution of revisions.

3.4.3 Signal precisions for AR(2)

Next I vary the precisions of the current and revised signals to understand their relationship to the bias caused by higher-order dynamics. Specifically, I normalize both precisions by the variance of the fundamental and vary the ratio $\frac{\tau_\nu}{\tau_\psi}$. When this ratio is greater than 1 the revised signal is more precise than the current signal. In the benchmark case this ratio was 4.

Table 3.7 simulates an AR(2) when τ_ψ ranges from 20% to 1000% the size of τ_ν. The bias caused by higher-order dynamics begins small and negative and grows more negative as τ_ψ rises. There is an inflection point when τ_ψ is half the value of τ_ν, at which point the bias begins to slowly fall towards 0, which it reaches only once current signal precision is four times as large as revised signal precision. Even at its maximum size, this bias is small, so the total bias is overwhelmingly driven by the revised signal. Total bias decreases as τ_ψ rises, since individuals shift weight from the revised signal to the relatively more accurate current signal.

3.4.4 Using calibrated data series

To give a rough idea of how strong the bias can be for real data series, I perform simulations in table 3.8 using the estimated parameters from section 3.2 and decompose both types of bias. As shown in the previous section, the size of the total bias heavily depends on the signal precisions, but these are inherently unobservable. Thus the results of this exercise should not be understood as true calibrated estimates of bias from measuring information frictions on

Table 3.7. Bias as signal precisions vary

$\frac{\tau_\psi}{\tau_v}$	**True K_1**	**Bias** **Total**	**Revisions**	**Dynamics**
0.2	0.22	-0.198	-0.194	-0.003
0.3	0.30	-0.198	-0.192	-0.006
0.4	0.36	-0.188	-0.180	-0.007
0.5	0.41	-0.174	-0.167	-0.007
0.6	0.45	-0.161	-0.154	-0.007
0.7	0.48	-0.148	-0.142	-0.006
0.8	0.51	-0.137	-0.131	-0.006
0.9	0.54	-0.127	-0.121	-0.005
1.0	0.57	-0.117	-0.112	-0.005
1.1	0.59	-0.109	-0.104	-0.005
1.2	0.61	-0.101	-0.097	-0.004
1.5	0.65	-0.083	-0.080	-0.003
2.0	0.71	-0.062	-0.060	-0.002
4.0	0.82	-0.03	-0.03	0.000
10.0	0.92	-0.01	-0.01	0.000

Notes: Total bias is the difference between the true K_1 and the estimated $\hat{K}_1$ produced by the original CG specification (3.2). Bias from revisions is the normalized covariance between the average forecast update and the fundamental shock w_t multiplied by $\left(\frac{K_r}{\rho_1 K_1 + K_r}\right)$. Bias from dynamics is the sum of normalized covariances between the average forecast update and lagged backcast errors multiplied by $\left(\frac{\rho_s K_r}{\rho_1 K_1 + K_r}\right)$ for $s \in (2, \ldots, p)$. Each row had 500 economies simulated, with 1000 individuals and 1000 time periods. The bias figures are averages across all simulations.

real data, but instead background understanding for how the bias may vary across different data series.[12]

Table 3.8. Calibrated simulations

	True K_1	Bias		
		Total	**Revisions**	**Dynamics**
Nominal GDP growth	0.35	-0.33	-0.33	0.00
Real GDP growth	0.35	-0.32	-0.32	0.00
GDP price inflation	0.36	-0.33	-0.33	0.01
Real consumption growth	0.34	-0.50	-0.50	0.00
Real nonresidential investment growth	0.35	-0.34	-0.34	0.00
Real residential investment growth	0.35	-0.28	-0.28	0.00
CPI inflation	0.36	-0.29	-0.29	0.00
Housing starts growth	0.34	-0.36	-0.36	0.00
Industrial production growth	0.36	-0.26	-0.26	-0.01
Unemployment rate	0.47	-0.12	-0.12	-0.03
Employment growth	0.37	-0.23	-0.23	-0.01
Treasury bill 3 month rate	0.42	-0.16	-0.16	-0.02
Treasury bond 10 year rate	0.40	-0.18	-0.18	-0.01
AAA Corporate bond rate	0.42	-0.17	-0.17	-0.01

Notes: The specific AR parameters for each variable are taken from table 3.1. Each row had 100 economies simulated, with 500 individuals and 500 time periods. The bias figures are averages across all simulations.

Most of the calibrated variables do not suffer bias due to higher-order dynamics, even strongly persistent series like CPI inflation or nominal GDP. Unemployment has the highest share of total bias caused by higher-order dynamics, of nearly 25%, due to its highly cyclical nature from complex conjugate roots. The other data series with relatively large bias contributions due to higher-order dynamics also have similar time series properties, such as the financial variables and employment growth. This is not surprising given the simplicity of the calibration exercise, but it helps demonstrate how the simulation results can apply to real-world data.

[12] Also note that many of these data series do not receive data revisions in the traditional sense, or any revisions at all (such as the CPI). As such the results in table 3.8 should be viewed as upper-bounds on the total bias.

3.4.5 Discussion

The main lesson from these simulations is that the bias caused by higher-order dynamics is small across nearly all parameterizations. This should come as no surprise given where it comes from: the covariance between the forecast update and lagged backcast errors. All of the non-zero covariance between these terms is based on values far into the past, starting with x_{t-2}, and as such they are heavily discounted by the Kalman filter mechanism. Combined with a generally low value for the Kalman gain on revised data, K_r, the magnitude of the bias resulting from the revised signal interacting with higher-order dynamics must necessarily be small.

More relevant is how the bias due to the revised signal alone may vary depending on the higher-order properties of the fundamental. Adding additional lags tended to lower the total bias, as did increasing the unconditional variance of the process. Holding that variance constant, the total bias was monotonically decreasing in the degree of first-order persistence, ρ_1. This suggests the interaction of information flows with time series behavior can have first-order effects on measurement of information frictions without considering a dynamic information structure.

3.5 Conclusion

I showed the leading method for estimating information frictions from surveys of forecasters can be applied without modification to variables with higher-order dynamics. This emphasizes the key mechanism of the CG method: using information flows to predict forecast errors. When the information framework changes, such as when it includes a revised signal, then higher-order dynamics introduce a novel bias in the CG method by interacting with observations of past values. Through numeric simulations I demonstrated that this bias is quantitatively small relative to the bias from observing revised data, which underlines the importance of information flows rather than higher-order dynamics. I proposed a correction term to the regression that isolates the additional information gained from the revised signal and verified its accuracy relative to the uncorrected CG method. These results suggest the corrected regression is the appropriate way to measure information frictions when individuals observe revised data, regardless of the time series properties of the variable they are forecasting.

Chapter 3 is currently being prepared for submission for publication; Tyler L. Paul. The dissertation author was the principal author on this paper.

Bibliography

Amador, Manuel, and Pierre-Olivier Weill. 2012. "Learning from Private and Public Observations of Others Actions." *Journal of Economic Theory*, 147(3): 910–940.

Angeletos, George-Marios, and Alessandro Pavan. 2004. "Transparency of Information and Coordination in Economies with Investment Complementarities." *American Economic Review*, 94(2): 91–98.

Angeletos, George-Marios, and Alessandro Pavan. 2007. "Efficient Use of Information and Social Value of Information." *Econometrica*, 75(4): 1103–1142.

Angeletos, George-Marios, and Chen Lian. 2018. "Forward Guidance without Common Knowledge." *American Economic Review*, 108(9): 2477–2512.

Angeletos, George-Marios, and Chen Lian. 2022. "Dampening General Equilibrium: Incomplete Information and Bounded Rationality." National Bureau of Economic Research Working Paper 29776.

Angeletos, George-Marios, and Jennifer La'O. 2020. "Optimal Monetary Policy with Informational Frictions." *Journal of Political Economy*, 128(3): 1027–1064.

Angeletos, George-Marios, and Zhen Huo. 2021. "Myopia and Anchoring." *American Economic Review*, 111(4): 1166–1200.

Aruoba, S. Boragan. 2008. "Data Revisions Are Not Well Behaved." *Journal of Money, Credit and Banking*, 40(2-3): 319–340.

Bańbura, Marta, Domenico Giannone, Michele Modugno, and Lucrezia Reichlin. 2013. "Chapter 4 - Now-Casting and the Real-Time Data Flow." In *Handbook of Economic Forecasting*. Vol. 2 of *Handbook of Economic Forecasting*, , ed. Graham Elliott and Allan Timmermann, 195–237. Elsevier.

Berger, Helge, Michael Ehrmann, and Marcel Fratzscher. 2011. "Geography, skills or both: What explains Fed watchers' forecast accuracy of US monetary policy?" *Journal of Macroeconomics*, 33(3): 420–437.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. 2020. "Overreaction in Macroeconomic Expectations." *American Economic Review*, 110(9): 2748–2782.

Box, George E.P., Gwilym M. Jenkins, and Gregory C. Reinsel. 1994. *Time Series Analysis: Forecasting and Control.* . Third ed., Englewood Cliffs, N.J.:Prentice Hall.

Candia, Bernardo, Michael Weber, Yuriy Gorodnichenko, and Olivier Coibion. 2023. "Perceived and Expected Rates of Inflation of US Firms." *AEA Papers and Proceedings*, 113: 52–55.

Charles, Amélie, and Olivier Darné. 2012. "Trends and random walks in macroeconomic time series: A reappraisal." *Journal of Macroeconomics*, 34(1): 167–180. Has macro progressed?

Clements, Michael P. 2012. "Do professional forecasters pay attention to data releases?" *International Journal of Forecasting*, 28(2): 297–308.

Clements, Michael P, and Ana Beatriz Galvão. 2017. "Data Revisions and Real-time Probabilistic Forecasting of Macroeconomic Variables." Henley Business School, University of Reading ICMA Centre Discussion Papers in Finance icma-dp2017-01.

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

Coibion, Olivier, and Yuriy Gorodnichenko. 2015. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review*, 105(8): 2644–2678.

Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko, and Michael Weber. 2021*a*. "Forward Guidance and Household Expectations." Department of Economics, Institute for Business and Economic Research, UC Berkeley Department of Economics, Working Paper Series qt7np4511v.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber. 2022. "Monetary Policy Communications and Their Effects on Household Inflation Expectations." *Journal of Political Economy*, 130(6): 1537–1584.

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. 2020. "Inflation Expectations and Firm Decisions: New Causal Evidence." *The Quarterly Journal of Economics*, 135(1): 165–219.

Coibion, Olivier, Yuriy Gorodnichenko, Saten Kumar, and Jane Ryngaert. 2021*b*. "Do You

Know that I Know that You Know…? Higher-Order Beliefs in Survey Data." *The Quarterly Journal of Economics*, 136(3): 1387–1446.

Cole, Rosanne. 1969. *Errors in Provisional Estimates of Gross National Product. NBER Books*, National Bureau of Economic Research, Inc.

Croushore, Dean, and Tom Stark. 2001. "A real-time data set for macroeconomists." *Journal of Econometrics*, 105(1): 111–130.

Dickey, David A., and Wayne A. Fuller. 1979. "Distribution of the Estimators for Autoregressive Time Series With a Unit Root." *Journal of the American Statistical Association*, 74(366): 427–431.

D'Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber. 2021. "Exposure to Grocery Prices and Inflation Expectations." *Journal of Political Economy*, 129(5): 1615–1639.

Gali, Jordi, Mark Gertler, and J. David Lopez-Salido. 2005. "Robustness of the estimates of the hybrid New Keynesian Phillips curve." *Journal of Monetary Economics*, 52(6): 1107–1118.

Galvão, Ana Beatriz. 2017. "Data revisions and DSGE models." *Journal of Econometrics*, 196(1): 215–232.

Gemmi, Luca, and Rosen Valchev. 2022. "Biased Surveys." *Working Paper*.

Goldstein, Nathan. 2019. "Tracking inattention." *Available at SSRN 3641353*.

Goldstein, Nathan, and Yuriy Gorodnichenko. 2022. "Expectations Formation and Forward Information." National Bureau of Economic Research, Inc NBER Working Papers 29711.

Hamilton, James Douglas. 1994. *Time series analysis.* Princeton, NJ:Princeton Univ. Press.

Hellwig, Christian. 2010. "Monetary Business Cycles (Imperfect Information)." *Monetary Economics*, , ed. Steven N. Durlauf and Lawrence E. Blume, 181–188. London:Palgrave Macmillan UK.

Holdren, Alyssa E. 2014. "Gross domestic product and gross domestic income: revisions and source data." *Survey of Current Business*, 94(6).

Howrey, E Philip. 1978. "The Use of Preliminary Data in Econometric Forecasting." *The Review of Economics and Statistics*, 60(2): 193–200.

Howrey, E Philip. 1984. "Data Revision, Reconstruction, and Prediction: An Application to Inventory Investment." *The Review of Economics and Statistics*, 66(3): 386–393.

Huo, Zhen, and Marcelo Pedroni. 2020. “A Single-Judge Solution to Beauty Contests.” *American Economic Review*, 110(2): 526–568.

Huo, Zhen, and Marcelo Pedroni. 2023. “Dynamic information aggregation: Learning from the past.” *Journal of Monetary Economics*, 136: 107–124.

Jacobs, Jan P.A.M., and Simon van Norden. 2011. “Modeling data revisions: Measurement error and dynamics of "true" values.” *Journal of Econometrics*, 161(2): 101–109.

Kamber, Güneş, James Morley, and Benjamin Wong. 2018. “Intuitive and Reliable Estimates of the Output Gap from a Beveridge-Nelson Filter.” *The Review of Economics and Statistics*, 100(3): 550–566.

Kasa, Kenneth. 2000. “Forecasting the Forecasts of Others in the Frequency Domain.” *Review of Economic Dynamics*, 3(4): 726–756.

Kishor, N. Kundan, and Evan F. Koenig. 2009. “VAR Estimation and Forecasting When Data Are Subject to Revision.” *Journal of Business & Economic Statistics*, 30(2): 181–190.

Kohlhas, Alexandre, and Tobias Broer. 2019. “Forecaster (Mis-)Behavior.” Society for Economic Dynamics 2019 Meeting Papers 1171.

Kohlhas, Alexandre N., and Ansgar Walther. 2021. “Asymmetric Attention.” *American Economic Review*, 111(9): 2879–2925.

Kumar, Saten, Hassan Afrouzi, Olivier Coibion, and Yuriy Gorodnichenko. 2015. “Inflation Targeting Does Not Anchor Inflation Expectations: Evidence from Firms in New Zealand.” *Brookings Papers on Economic Activity*, 46(2 (Fall)): 151–225.

Lucas, Robert Jr. 1972. “Expectations and the neutrality of money.” *Journal of Economic Theory*, 4(2): 103–124.

Mankiw, N. Gregory, and Matthew D. Shapiro. 1986. “News or Noise? An Analysis of GNP Revisions.” National Bureau of Economic Research, Inc NBER Working Papers 1939.

Mankiw, N. Gregory, and Ricardo Reis. 2002. “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve.” *The Quarterly Journal of Economics*, 117(4): 1295–1328.

Mankiw, N Gregory, David E Runkle, and Matthew D Shapiro. 1984. “Are preliminary announcements of the money stock rational forecasts?” *Journal of Monetary Economics*, 14(1): 15–27.

McClure, Ethan M. L, Olivier Coibion, and Yuriy Gorodnichenko. 2022. “The Macroeconomic Expectations of U.S. Managers." National Bureau of Economic Research Working Paper 29986.

Melosi, Leonardo. 2016. "Signalling Effects of Monetary Policy." *The Review of Economic Studies*.

Mester, Loretta J. 2016. "Acknowledging Uncertainty, 10-07-2016; Shadow Open Market Committee Fall Meeting, New York, NY." Federal Reserve Bank of Cleveland Speech 77.

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

Muth, John F. 1960. "Optimal properties of exponentially weighted forecasts." *Journal of the American Statistical Association*, 55(290): 299–306.

Nimark, Kristoffer P. 2014. "Man-Bites-Dog Business Cycles." *American Economic Review*, 104(8): 2320–2367.

Pesaran, M. Hashem, and Martin Weale. 2006. "Survey Expectations." In *Handbook of Economic Forecasting*. Vol. 1 of *Handbook of Economic Forecasting*, , ed. G. Elliott, C. Granger and A. Timmermann, Chapter 14, 715–776. Elsevier.

Rondina, Giacomo. 2008. "Incomplete Information and Informative Pricing: Theory and Application." Society for Economic Dynamics 2008 Meeting Papers 981.

Ropele, Tiziano, Yuriy Gorodnichenko, and Olivier Coibion. 2024. "Inflation Expectations and Misallocation of Resources: Evidence from Italy." *American Economic Review: Insights*, 6(2): 246–61.

Ryngaert, Jane. 2017. "What Do (and Don't) Forecasters Know About US Inflation?" *University of Texas at Austin*.

Sargent, Thomas J. 1989. "Two Models of Measurements and the Investment Accelerator." *Journal of Political Economy*, 97(2): 251–287.

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

Stark, Tom. 2013. "SPF panelists' forecasting methods: A note on the aggregate results of a November 2009 special survey." *Federal Reserve Bank of Philadelphia*.

Stark, Tom, and Dean Croushore. 2002. "Forecasting with a real-time data set for macroeconomists." *Journal of Macroeconomics*, 24(4): 507–531.

Svensson, Lars E. O. 2006. "Social Value of Public Information: Comment: Morris and Shin (2002) Is Actually Pro-Transparency, Not Con." *American Economic Review*, 96(1): 448–452.

Townsend, Robert. 1983. "Forecasting the Forecasts of Others." *Journal of Political Economy*, 91(4): 546–88.

U.S. Bureau of Economic Analysis. 2024. "Vintage History of Quarterly Gross Domestic Product (GDP) and Gross Domestic Income (GDI) Estimates." U.S. Bureau of Economic Analysis.

U.S. Bureau of Labor Statistics. 2024. "Household and establishment survey response rates." U.S. Bureao of Labor Statistics.

U.S. Department of Labor. 2024. "Congressional Budget Justification Bureau of Labor Statistics." U.S. Department of Labor.

Woodford, Michael. 2001. "Imperfect Common Knowledge and the Effects of Monetary Policy." National Bureau of Economic Research, Inc NBER Working Papers 8673.

A.1 Appendix - Chapter 1

A.1.1 State-space representation

I write the information framework in the following state-space representation:

$$\mathbf{\xi_t} = \mathbf{F}\mathbf{\xi_{t-1}} + \mathbf{w_t}$$

$$\mathbf{Z_{it}} = \mathbf{H}'\mathbf{\xi_t} + \mathbf{v_t}$$

$$\mathbf{\xi_t} = \begin{bmatrix} x_t \\ x_{t-1} \end{bmatrix} \quad \mathbf{F} = \begin{bmatrix} \rho & 0 \\ 1 & 0 \end{bmatrix} \quad \mathbf{w_t} = \begin{bmatrix} w_t \\ 0 \end{bmatrix}$$

$$\mathbf{H}' = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \mathbf{v_t} = \begin{bmatrix} \psi_{it} \\ \nu_t \end{bmatrix}$$

where $\mathbf{Q} = \mathbb{E}[\mathbf{w_t}\mathbf{w_t'}] = \begin{bmatrix} \tau_w^{-1} & 0 \\ 0 & 0 \end{bmatrix}$ and $\mathbf{R} = \mathbb{E}[\mathbf{v_t}\mathbf{v_t'}] = \begin{bmatrix} \tau_\psi^{-1} & 0 \\ 0 & \tau_\nu^{-1} \end{bmatrix}$. Applying the Kalman filter provides the recursive forecasting formula:

$$\mathbb{E}_{it}[\boldsymbol{\xi}_t] = \boldsymbol{F}\mathbb{E}_{it-1}[\boldsymbol{\xi}_{t-1}] + \boldsymbol{K}(\boldsymbol{Z}_{it} - \boldsymbol{H'}\boldsymbol{F}\mathbb{E}_{it-1}[\boldsymbol{\xi}_{t-1}])$$

which gives us equation (1.1) in the text. The matrix $\boldsymbol{K}$ is the Kalman gain matrix. The first row corresponds to weights placed on the signals for forecasting x_t, and the second row corresponds to weights placed on the signals for for forecasting x_{t-1}. The matrix is a function of the squared forecast error matrix is a function of the squared forecast error matrix, $\boldsymbol{P}_{t+1} = \mathbb{E}[\mathbb{E}_{it}[\boldsymbol{\xi}_{t+1}] - \boldsymbol{\xi}_{t+1}]^2$, which I find by solving the following Ricatti equation:

$$\boldsymbol{P}_{t+1} = \boldsymbol{F}(\boldsymbol{P}_{t+1} - (\boldsymbol{P}_{t+1}\boldsymbol{H}(\boldsymbol{H'}\boldsymbol{P}_{t+1}\boldsymbol{H} + \boldsymbol{R})^{-1}\boldsymbol{H'}\boldsymbol{P}_{t+1}))\boldsymbol{F'} + \boldsymbol{Q}$$

Once solved, I can write each gain parameter as follows:

$$
\begin{aligned}
K &= \frac{(\tau_\psi)(1+\kappa(\rho^2\tau_w+\tau_v))}{(\tau_\psi)(1+\kappa(\rho^2\tau_w+\tau_v))+\tau_w(1+\tau_v\kappa)} \\
K_r &= \frac{\rho\tau_w\tau_v\kappa}{(\tau_\psi)(1+\kappa(\rho^2\tau_w+\tau_v))+\tau_w(1+\tau_v\kappa)} \\
K^b &= \frac{\rho\tau_w(\tau_\psi)\kappa}{(\tau_\psi)(1+\kappa(\rho^2\tau_w+\tau_v))+\tau_w(1+\tau_v\kappa)} \\
K_r^b &= \frac{\tau_v\kappa(\tau_w+\tau_\psi)}{(\tau_\psi)(1+\kappa(\rho^2\tau_w+\tau_v))+\tau_w(1+\tau_v\kappa)} \\
\kappa &= \frac{1+\kappa(\rho^2\tau_w+\tau_v)}{(\tau_\psi)(1+\kappa(\rho^2\tau_w+\tau_v))+\tau_w(1+\tau_v\kappa)}
\end{aligned}
$$

where κ is the filter's steady-state mean-squared error, and K^b refers to the Kalman gain parameters applied to both signals when forecasting x_{t-1}.

A.1.2 Modeling revisions as news

When a revision contains news it must reveal information about the state that was unavailable in the previous period. To model this I assume the fundamental is composed of two shocks, only one of which is noisily observed in the initial period. For ease of exposition I assume the revised signal is also idiosyncratic. the new model is:

$$\text{State: } x_t = \rho x_{t-1} + u_t + w_t$$
$$\text{Current signal: } y_{it} = \rho x_{t-1} + u_t + \psi_{it}$$
$$\text{Revised signal: } r_t = x_{t-1} + \nu_{it}$$

The current signal is a noisy observation of the state less the fundamental shock w_t. This shock is only observed in the revised signal one period later. The optimal forecasting equation when $h = 0$ is:

$$\begin{aligned}\mathbb{E}_{it}[x_t] &= \mathbb{E}_{it-1}[x_t] + K(y_{it} - \mathbb{E}_{it-1}[y_{it}]) + K_r(r_t - \mathbb{E}_{it-1}[r_t]) \\ &= \left(\frac{\rho - \rho K - K_r}{\rho}\right)\mathbb{E}_{it-1}[x_t] + K y_{it} + K_r r_t\end{aligned}$$

To convert the above into a relationship between forecast error and forecast update, I use the following two equivalences:

$$\begin{aligned}y_{it} &= \rho x_{t-1} + u_t + \psi_{it} = x_t - u_t + \psi_{it} \\ r_t &= \left(\frac{1}{\rho}\right)(x_t - w_t - u_t) + \nu_{it}\end{aligned}$$

Similar to section 1.2.2, in order to convert both signals into equivalent information content about x_t I must account for the differences in fundamental shocks observed. Once averaged across forecasters and rearranged, the relationship between forecast error and forecast update is:

$$x_t - \mathbb{E}_t[x_t] = \left(\frac{\rho - \rho K - K_r}{\rho K + K_r} \right) (\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t]) + u_t + \left(\frac{K_r w_t}{\rho K + K_r} \right)$$

This relationship is similar to equation (1.4), with the fundamental shock u_t and w_t both appearing in the residual. Thus the bias when estimating equation (1.3) is still present when data revisions contain news rather than noise.

A.1.3 Additional revised signals

Since many data series undergo several rounds of revisions, I now consider a model with two revised signals. The new model is:

$$\text{State: } x_t = \rho x_{t-1} + w_t$$

$$\text{Current signal: } y_{it} = x_t + \psi_{it}$$

$$\text{1st Revised signal: } r_{1t} = x_{t-1} + \nu_{1t}$$

$$\text{2nd Revised signal: } r_{2t} = x_{t-2} + \nu_{2t}$$

The optimal forecasting equation is:

$$\mathbb{E}_{it}[x_t] = \mathbb{E}_{it-1}[x_t] + K(y_{it} - \mathbb{E}_{it-1}[x_t]) + K_{1r}(r_{1t} - \mathbb{E}_{it-1}[x_{t-1}]) + K_{2r}(r_{2t} - \mathbb{E}_{it-1}[x_{t-2}]) \quad \text{(A.11)}$$

which now includes the difference between the second revised signal and the backcast of x_{t-2} made in period $t-1$. With similar manipulations to the main text I find the following relationship between average forecast error and average forecast update:

$$x_t - \mathbb{E}_t[x_t] = \left(\frac{\rho^2 - \Lambda}{\Lambda}\right)(\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t]) + \left(\frac{\rho K_{2r}}{\Lambda}\right)(w_{t-1} - \mathbb{E}_{t-1}[w_{t-1}]) \quad \text{(A.12)}$$

$$+ \left(\frac{\rho K_{1r} + K_{2r}}{\Lambda}\right) w_t - \left(\frac{\rho^2}{\Lambda}\right) \varepsilon_t \quad \text{(A.13)}$$

where $\Lambda = \rho^2 K + \rho K_{1r} + K_{2r}$ and $\varepsilon_t = K_{1r}\nu_{1t} + K_{2r}\nu_{2t}$. In addition to the same upward bias caused by the current period fundamental w_t, there is a new term which entirely depends on the presence of the second revised signal (in other words, when $K_{2r} \neq 0$). This measures the forecast error of predicting w_{t-1} given information in period $t-1$. Intuitively, this term appears because the forecaster made a (non-zero) prediction of w_{t-1} in period $t-1$. Since the difference in information content between r_{1t} and r_{2t} is w_{t-1}, to account for the new information I must include the forecast error of w_{t-1}. The mechanism is identical to the appearance of w_t as described in the main text; however, because the forecast of w_t in period $t-1$ is the unconditional mean of 0, there was no expectations term associated with the w_t term.

If I apply my correction term I find the following relationship:

$$\begin{aligned} x_t - \mathbb{E}_t[x_t] = \left(\frac{\rho^2 - \lambda}{\lambda}\right)(\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t]) - \frac{\rho^2 K_{1r}}{\lambda}(r_{1t} - \mathbb{E}_{t-1}[x_{t-1}]) \\ + \left(\frac{\rho K_{2r}}{\lambda}\right)(w_{t-1} - \mathbb{E}_{t-1}[w_{t-1}]) + \left(\frac{K_{2r}}{\lambda}\right) w_t - \left(\frac{K_{2r}}{\lambda}\right) \nu_{2t} \end{aligned}$$

where $\lambda = \rho^2 K + K_{2r}$. By controlling for the impact of r_{1t} on the forecast update, I remove K_{1r} from Λ as well as the weight placed on w_t and the noise term ν_{1t}. However, the correction term is now insufficient to make the residual orthogonal to the regressors. The fundamental shock w_t is still present in addition to last period's fundamental forecast error. The latter is correlated with both the forecast update and the revised signal innovation, while the former is only correlated with the forecast update.

There are two takeaways from this exercise. First, if the true information structure available to forecasters includes multiple revised signals, my regression specification (1.17) will produce biased estimates of both β and β_r. These biases would confound the estimates in section 1.4 which may partially explain some of the null findings. On the other hand, the size of the bias depends on K_{2r} which may be small in practice.

Second, each additional revised signal introduces a new 'bump' to ones forecast, as past expectations of fundamental shocks are revised according to the new information. In a generalized environment with m revised signals we'd expect relationship (A.13) to have m terms: $w_t, w_{t-1}, \ldots, w_{t-m}$. This is the case when the state follows an AR(1). I anticipate the interaction with p lags of the state to be more complex.

A.1.4 Modeling NIPA data releases with a lag

Since NIPA variables are published with a full quarter lag, this section considers an alternative model to capture these dynamics. Now I let $y_{it} = x_{t-1} + \psi_{it}$ and $r_t = x_{t-2} + \nu_t$ be the signals. After performing the same manipulations as before, I find a slightly different relationship than (1.16):

$$x_{t+h} - \mathbb{E}_t[x_{t+h}] = \left(\frac{\rho - K}{K}\right)(\mathbb{E}_t[x_{t+h}] - \mathbb{E}_{t-1}[x_{t+h}]) - \left(\frac{\rho^{h+1} K_r}{K}\right)(r_t - \mathbb{E}_{t-1}[x_{t-2}]) + \rho\tilde{\gamma}_t \tag{A.14}$$

where $\tilde{\gamma}_t = \sum_{j=0}^{h-1} \rho^{h-1-j} w_{t+j}$ is a summation of fundamental shocks beginning with w_t, and thus unobserved by forecasters in period t under this new informational assumption. To recover the degree of information frictions from (A.14), I need both an estimate of ρ and the average backcast $\mathbb{E}_{t-1}[x_{t-2}]$. As mentioned previously, the SPF has backcasts for the level of all variables, so this is empirically feasible.

A.1.5 Proofs of propositions

Proposition 1

First I will prove the difference equation representation of the average forecast error. Expanding out x_{t+h} and plugging in the average forecast equation, average forecast error looks like:

$$x_{t+h} - \mathbb{E}_t[x_{t+h}] = \rho^h x_t + \gamma_{t+h} - \mathbb{E}_{t-1}[x_{t+h}] - \rho^h K(x_t - \mathbb{E}_{t-1}[x_{t+h}])$$

where $\gamma_{t+h} = \sum_{j=1}^{h} \rho^{h-j} w_{t+j}$. Iterating x_t backwards one period and re-arranging gives the equation in the proposition.

$$\begin{aligned}
&= \rho^{h+1} x_{t-1} + \rho^h w_t + \gamma_{t+h} - \rho^{h+1}\mathbb{E}_{t-1}[x_{t-1}] - \rho^h K(\rho x_{t-1} + w_t - \rho^{h+1}\mathbb{E}_{t-1}[x_{t-1}]) \\
&= \rho^{h+1}(1-K)(x_{t-1} - \mathbb{E}_{t-1}[x_{t-1}]) + \rho^h(1-K)w_t + \gamma_{t+h}
\end{aligned}$$

To find the average forecast update, subtract $\mathbb{E}_{t-1}[x_{t+h}]$ from both sides of the average forecast equation, iterate x_t backwards one period and rearrange.

Proposition 2

To find the structural relationship with an informative revised signal I start from the average forecasting equation. Using the result from the text, I can rewrite it as follows:

$$\begin{aligned}
\mathbb{E}_t[x_{t+h}] &= \rho^h \mathbb{E}_{t-1}[x_t] + \rho^h K(x_t - \mathbb{E}_{t-1}[x_t]) + \rho^h K_r \frac{(x_t - \mathbb{E}_{t-1}[x_t]) - w_t + \rho \nu_t}{\rho} \\
\rho^h \mathbb{E}_t[x_t] &= \rho^h (1 - K - \frac{K_r}{\rho}) E_{t-1}[x_t] + \rho^h (K + \frac{K_r}{\rho}) x_t - \rho^h K_r \frac{(w_t - \rho \nu_t)}{\rho}
\end{aligned}$$

Cancel out ρ^h, subtract $\frac{\rho(1-K)-K_r}{\rho}\mathbb{E}_{t-1}[x_t]$ and $\frac{\rho K + K_r}{\rho}\mathbb{E}_t[x_t]$ from both sides to find:

$$\frac{\rho(1-K) - K_r}{\rho}(\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t]) = (K + \frac{K_r}{\rho})(x_t - \mathbb{E}_t[x_t]) - K_r \frac{(w_t - \rho \nu_t)}{\rho}$$

Rearranging terms and dividing through by $\frac{\rho K+K_r}{\rho}$ gives us the equation in proposition 2.

$$\frac{\rho(1-K)-K_r}{\rho K+K_r}(\mathbb{E}_t[x_t]-\mathbb{E}_{t-1}[x_t])+K_r\left(\frac{w_t-\rho \nu_t}{\rho K+K_r}\right)=x_t-\mathbb{E}_t[x_t]$$

Proposition 3

This section shows the steps to compute the bias equation (1.5). First I address the numerator:

$$\mathbb{E}\left[(\mathbb{E}_t[x_{t+h}]-\mathbb{E}_{t-1}[x_{t+h}])\left(\left(\frac{\rho^h K_r}{\rho K+K_r}\right)w_t-\left(\frac{\rho^{h+1}K_r \nu_t}{\rho K+K_r}\right)\right)\right]$$

where I have ignored the term γ_{t+h} as argued in the body of the text. Using the forecasting equation I can rewrite the forecast update as:

$$\mathbb{E}_t[x_{t+h}]-\mathbb{E}_{t-1}[x_{t+h}]=\rho^h\left(Kw_t+(\rho K+K_r)x_{t-1}+K_r\nu_t-(\rho K+K_r)\mathbb{E}_{t-1}[x_{t-1}]\right)$$

The two values in the error term $\varepsilon_t=\rho^h\frac{K_r w_t-\rho K_r \nu_t}{\rho K+K_r}$, w_t and ν_t, are i.i.d. and dated time t, so there is no correlation with expectations or values from $t-1$ or earlier. Multiplying together, I find the numerator of the bias is

$$(\mathbb{E}_t[x_{t+h}]-\mathbb{E}_{t-1}[x_{t+h}])\varepsilon_t=\rho^{2h}\left(\frac{KK_r\tau_w^{-1}-\rho K_r^2\tau_\nu^{-1}}{\rho K+K_r}\right)$$

To compute the denominator I must find the squared expectation of the average forecast update. To start I use the representation of the average forecast update from equation (1.13):

$$\mathbb{E}_t[x_{t+h}]-\mathbb{E}_{t-1}[x_{t+h}]=\rho^h\frac{(Kw_t+K_r w_{t-1})+(1-\rho L)(K_r\nu_t)}{1-GL}$$

where $G = (\rho - \rho K - \rho K_r)$. I represent the first term in w_t and w_{t-1} of this expression as the following ARMA(1,1):

$$\tilde{h}_t = G\tilde{h}_{t-1} + \tilde{w}_t + \frac{K_r}{K}\tilde{w}_{t-1}$$

where $\tilde{w}_t = \rho^h K w_t$. Let the second term be

$$\tilde{j}_t = G\tilde{j}_{t-1} + \tilde{v}_t - \rho\tilde{v}_{t-1}$$

where $\tilde{v}_t = K_r v_t$. Since $\tilde{w}_{t+s}$ and $\tilde{v}_{t+u}$ are independent of each other for all s and u, I can ignore the cross-product term when squaring and only consider the autocovariances of both $\tilde{h}_t$ and $\tilde{j}_t$. I find the autocovariances by solving the Yule-Walker equation:

$$\gamma_k = G\gamma_{k-1} + \mathbb{E}[\tilde{w}_t\tilde{h}_{t-k} + \frac{K_r}{K}\tilde{w}_{t-1}\tilde{h}_{t-k}]$$

When $k > 1$, the terms in the bracket equal 0. When $k = 1$, only the second term is non-zero:

$$\gamma_1 = G\gamma_0 + \mathbb{E}\left[\frac{K_r}{K}\tilde{w}_{t-1}\tilde{h}_{t-1}\right] = G\gamma_0 + \rho^{2h}KK_r\tau_w^{-1}$$

When $k = 0$ the equation becomes:

$$\gamma_0 = G\gamma_1 + \rho^{2h}\tau_w^{-1}(K^2 + K_r^2 + GKK_r) = G\gamma_1 + \rho^{2h}\tau_w^{-1}(K^2 + K_r(1-K)(\rho K + K_r))$$

Plugging in γ_1 and solving, I find the autocovariance of h_t is:

$$\mathbb{E}[h_t^2] = \gamma_0 = \rho^{2h}\tau_w^{-1}\frac{(K^2 + K_r^2 + GKK_r)}{1 - G^2}$$

Using the same method for j_t, I find its autocovariance is:

$$\mathbb{E}[j_t^2] = \gamma_0 = \rho^{2h}K_r^2\tau_v^{-1}\frac{(1 - 2\rho G + \rho^2)}{1 - G^2}$$

Summing both together gives the denominator of the bias term, which when combined with the numerator gives equation (1.5) in the text. Notice how the horizon enters the numerator and denominator equally, which causes it to cancel and not affect the size of the bias.

A.1.6 Predicting forecast updates with data revisions, other horizons and controls

Table A.9. Data Revisions Predicting Aggregate Forecast Updates, h=0

Forecast Updates	Data Revision	p-value
Nominal GDP	-0.13	0.35
Real GDP	0.01	0.95
GDP price index inflation	0.03	0.87
Real consumption	**0.15**	0.05
Real nonresidential investment	0.00	0.94
Real residential investment	0.01	0.95
Consumer price index	**0.55**	0.00
Housing starts	**0.23**	0.00
Industrial production	0.01	0.59
Unemployment rate	**1.21**	0.00
Employment growth	**0.60**	0.00
Treasury bill 3 month rate	**1.00**	0.01
Treasury bond 10 year rate	**2.00**	0.00
AAA Corporate bond rate	**1.13**	0.00

Notes: Each row shows the estimated coefficient $\hat{\alpha}_1$ from regression model (1.18) with $h = 0$. The variables in the top panel are from the NIPA. Data used for estimation are from 1968Q4 to 2019Q4. Standard errors are computed using the Newey-West method with automatic bandwidth selection to correct for serial correlation. Estimates are bold if I reject the null hypothesis of ignoring revised data releases at the 10% level.

Table A.10. Data Revisions Predicting Aggregate Forecast Updates, h=1

Forecast Updates	Data Revision	p-value
Nominal GDP	-0.05	0.58
Real GDP	**0.11**	0.05
GDP price index inflation	0.11	0.45
Real consumption	**0.11**	0.09
Real nonresidential investment	0.03	0.43
Real residential investment	0.02	0.63
Consumer price index	0.04	0.24
Housing starts	-0.05	0.19
Industrial production	**0.01**	0.02
Unemployment rate	**1.36**	0.00
Employment growth	**0.38**	0.00
Treasury bill 3 month rate	**1.32**	0.00
Treasury bond 10 year rate	**2.01**	0.00
AAA Corporate bond rate	**1.15**	0.00

Notes: Each row shows the estimated coefficient $\hat{\alpha}_1$ from regression model (1.18) with $h = 1$. The variables in the top panel are from the NIPA. Data used for estimation are from 1968Q4 to 2019Q4. Standard errors are computed using the Newey-West method with automatic bandwidth selection to correct for serial correlation. Estimates are bold if I reject the null hypothesis of ignoring revised data releases at the 10% level.

Table A.11. Data Revisions Predicting Aggregate Forecast Updates, h=2

Forecast Revisions	Data Revision	p-value
Nominal GDP	-0.04	0.67
Real GDP	-0.09	0.34
GDP price index inflation	**0.18**	0.08
Real consumption	0.00	1.00
Real nonresidential investment	0.04	0.20
Real residential investment	0.01	0.90
Consumer price index	**0.05**	0.01
Housing starts	**-0.08**	0.03
Industrial production	**0.01**	0.03
Unemployment rate	**1.40**	0.00
Employment growth	**0.23**	0.00
Treasury bill 3 month rate	**1.38**	0.00
Treasury bond 10 year rate	**1.90**	0.00
AAA Corporate bond rate	**1.11**	0.00

Notes: Each row shows the estimated coefficient $\hat{\alpha}_1$ from regression model (1.18) with $h = 2$. The variables in the top panel are from the NIPA. Data used for estimation are from 1968Q4 to 2019Q4. Standard errors are computed using the Newey-West method with automatic bandwidth selection to correct for serial correlation. Estimates are bold if I reject the null hypothesis of ignoring revised data releases at the 10% level.

Table A.12. Data Revisions Predicting Aggregate Forecast Updates, h=3

Forecast Revisions	**Data Revision**	**p-value**
Nominal GDP	0.08	0.21
Real GDP	0.11	0.33
GDP price index inflation	0.06	0.65
Real consumption	0.08	0.16
Real nonresidential investment	-0.01	0.60
Real residential investment	-0.03	0.26
Consumer price index	0.02	0.36
Housing starts	-0.02	0.23
Industrial production	**0.01**	0.02
Unemployment rate	**1.34**	0.00
Employment growth	**0.13**	0.00
Treasury bill 3 month rate	**1.41**	0.00
Treasury bond 10 year rate	**1.76**	0.00
AAA Corporate bond rate	**1.06**	0.00

Notes: Each row shows the estimated coefficient $\hat{\alpha}_1$ from regression model (1.18) with $h = 3$. The variables in the top panel are from the NIPA. Data used for estimation are from 1968Q4 to 2019Q4. Standard errors are computed using the Newey-West method with automatic bandwidth selection to correct for serial correlation. Estimates are bold if I reject the null hypothesis of ignoring revised data releases at the 10% level.

Table A.13. Data Revisions Predicting Aggregate Forecast Updates
With controls, one-year ahead

Forecast Updates	Data Revision	p-value	Realtime signal	p-value
Nominal GDP	-0.15	0.51	**0.17**	0.01
Real GDP	-0.02	0.93	**0.18**	0.01
GDP price index inflation	0.06	0.88	**0.28**	0.00
Real consumption	0.26	0.23	0.12	0.34
Real nonresidential investment	-0.08	0.51	**0.36**	0.00
Real residential investment	-0.05	0.81	0.17	0.17
Consumer price index	-0.05	0.18	**0.12**	0.00
Housing starts	-0.05	0.54	**0.19**	0.00
Industrial production	**0.03**	0.09	0.02	0.80
Unemployment rate	**1.31**	0.00	0.02	0.14
Employment growth	0.23	0.38	**0.75**	0.00
Treasury bill 3 month rate	**1.41**	0.00	0.00	0.99
Treasury bond 10 year rate	**1.70**	0.00	**0.04**	0.01
AAA Corporate bond rate	**1.13**	0.00	**0.03**	0.01

Notes: Each row shows the estimated coefficient $\hat{\alpha}_1$ from regression model (1.18) for one-year ahead forecasts, as well as the estimated coefficient on the first data release. The variables in the top panel are from the NIPA. Data used for estimation are from 1968Q4 to 2019Q4. Standard errors are computed using the Newey-West method with automatic bandwidth selection to correct for serial correlation. Estimates are bold if I reject the null hypothesis of ignoring revised data releases at the 10% level.

A.1.7 Main empirical results for all horizons

Table A.14. Average Forecast Errors and Average Forecast Updates, h=0

	CG Model			Revised Model				
Variable	β	**pval**	R^2	β	**pval**	β_r	**pval**	R^2
Nominal GDP	**0.43**	0.00	0.06	**0.43**	0.00	-0.12	0.10	0.06
Real GDP	**0.37**	0.09	0.04	**0.38**	0.05	-0.05	0.34	0.04
GDP price index inflation	0.42	0.11	0.06	**0.40**	0.09	0.04	0.71	0.05
Real consumption	0.13	0.45	0.00	0.07	0.72	-0.07	0.26	-0.01
Real nonresidential investment	**0.56**	0.00	0.04	**0.53**	0.01	0.05	0.59	0.04
Real residential investment	**0.73**	0.00	0.10	**0.77**	0.00	-0.06	0.27	0.09
CPI inflation	**0.80**	0.00	0.39	**1.08**	0.00	**-0.43**	0.00	0.53
Housing starts	**0.38**	0.00	0.05	**0.54**	0.00	**0.37**	0.00	0.17
Industrial production	**0.50**	0.00	0.15	**0.50**	0.00	0.05	0.54	0.15
Unemployment rate	**0.22**	0.00	0.15	**0.43**	0.00	**-0.46**	0.00	0.23
Employment growth	**0.54**	0.00	0.40	**0.53**	0.00	**0.07**	0.08	0.39
Treasury bill 3 month rate	**0.22**	0.00	0.20	**0.16**	0.00	**0.37**	0.00	0.32
Treasury bond 10 year rate	**0.15**	0.00	0.09	**0.15**	0.00	-0.13	0.36	0.10
AAA Corporate bond rate	**0.10**	0.02	0.02	-0.08	0.32	**0.49**	0.00	0.18

Notes: The CG model refers to estimates of (1.3), while the Revised Model refers to estimates of (1.17). All forecasts are of x_t. The variables in the top panel are from the NIPA. Data used for estimation are from 1968Q4 to 2019Q4. Standard errors are computed using the Newey-West method with automatic bandwidth selection to correct for serial correlation. Estimates that are in bold are statistically significant at the 10% level.

Table A.15. Comparing estimates of K between CG and Revised model, h=0

Variable	Test stat	pval	CG K	REV K	Difference
Nominal GDP	0.07	0.79	0.70	0.70	0.00
Real GDP	1.02	0.31	0.73	0.72	0.00
GDP price index inflation	2.53	0.11	0.71	0.71	-0.01
Real consumption	3.00	0.08	0.88	0.93	**-0.05**
Real nonresidential investment	0.16	0.69	0.64	0.65	-0.01
Real residential investment	1.02	0.31	0.58	0.57	0.01
CPI inflation	12.65	0.00	0.56	0.48	**0.07**
Housing starts	2.84	0.09	0.72	0.65	**0.07**
Industrial production	2.09	0.15	0.66	0.67	0.00
Unemployment rate	16.76	0.00	0.82	0.70	**0.12**
Employment growth	0.96	0.33	0.65	0.65	-0.01
Treasury bill 3 month rate	3.44	0.06	0.82	0.86	**-0.04**
Treasury bond 10 year rate	0.05	0.82	0.87	0.87	0.00
AAA Corporate bond rate	11.16	0.00	0.91	1.09	**-0.18**

Notes: The test statistic is a chi-squared adjusted for heteroskedasticity. CG K and REV K refer to the K recovered from both model estimates according to the expression $K = \frac{1}{1+\hat{\beta}}$. Estimates that are in bold are statistically different at the 10% level.

Table A.16. Average Forecast Errors and Average Forecast Updates, h=1

	CG Model			Revised Model				
Variable	β	**pval**	R^2	β	**pval**	β_r	**pval**	R^2
Nominal GDP	**0.83**	0.00	0.05	**0.85**	0.00	-0.11	0.29	0.05
Real GDP	**0.77**	0.02	0.06	**0.82**	0.02	-0.14	0.27	0.06
GDP price index inflation	**0.76**	0.09	0.07	0.67	0.13	**0.21**	0.02	0.09
Real consumption	-0.01	0.95	-0.01	0.08	0.81	**0.13**	0.06	0.00
Real nonresidential investment	**0.88**	0.02	0.05	**0.88**	0.06	0.10	0.35	0.05
Real residential investment	**1.05**	0.00	0.07	**1.04**	0.00	**0.15**	0.07	0.08
CPI inflation	0.52	0.10	-0.01	**0.82**	0.00	**-0.24**	0.01	0.01
Housing starts	**0.82**	0.01	0.02	**1.03**	0.01	0.14	0.21	0.03
Industrial production	**0.91**	0.01	0.08	**0.91**	0.01	**-0.16**	0.05	0.09
Unemployment rate	**0.42**	0.00	0.12	**0.62**	0.00	**-0.53**	0.02	0.15
Employment growth	**1.50**	0.00	0.36	**1.39**	0.00	**0.31**	0.00	0.38
Treasury bill 3 month rate	**0.27**	0.00	0.07	**0.20**	0.01	**0.39**	0.06	0.08
Treasury bond 10 year rate	0.04	0.68	-0.01	0.05	0.64	-0.17	0.53	-0.01
AAA Corporate bond rate	0.10	0.12	0.00	-0.09	0.44	**0.56**	0.03	0.05

Notes: The CG model refers to estimates of (1.3), while the Revised Model refers to estimates of (1.17). All forecasts are of x_{t+1}. The variables in the top panel are from the NIPA. Data used for estimation are from 1968Q4 to 2019Q4. Standard errors are computed using the Newey-West method with automatic bandwidth selection to correct for serial correlation. Estimates that are in bold are statistically significant at the 10% level.

Table A.17. Comparing estimates of *K* between CG and Revised model, h=1

Variable	Test stat	pval	CG K	REV K	Difference
Nominal GDP	0.85	0.36	0.55	0.54	0.01
Real GDP	1.57	0.21	0.56	0.55	0.01
GDP price index inflation	1.15	0.28	0.57	0.60	-0.03
Real consumption	1.81	0.18	1.01	0.93	0.09
Real nonresidential investment	0.00	1.00	0.53	0.53	0.00
Real residential investment	0.00	0.97	0.49	0.49	0.00
CPI inflation	1.31	0.25	0.66	0.55	0.11
Housing starts	1.28	0.26	0.55	0.49	0.06
Industrial production	0.00	0.96	0.52	0.52	0.00
Unemployment rate	4.48	0.03	0.70	0.62	**0.09**
Employment growth	3.66	0.06	0.40	0.42	**-0.02**
Treasury bill 3 month rate	2.28	0.13	0.79	0.83	-0.04
Treasury bond 10 year rate	1.63	0.20	0.96	0.95	0.01
AAA Corporate bond rate	4.08	0.04	0.91	1.10	**-0.18**

Notes: The test statistic is a chi-squared adjusted for heteroskedasticity. CG K and REV K refer to the K recovered from both model estimates according to the expression $K = \frac{1}{1+\hat{\beta}}$. Estimates that are in bold are statistically different at the 10% level.

Table A.18. Average Forecast Errors and Average Forecast Updates, h=2

Variable	CG Model			Revised Model				
	β	pval	R^2	β	pval	β_r	pval	R^2
Nominal GDP	0.40	0.25	0.00	0.39	0.27	0.09	0.19	0.00
Real GDP	**0.87**	0.02	0.05	**0.86**	0.02	0.07	0.45	0.04
GDP price index inflation	**1.10**	0.01	0.08	**0.93**	0.01	**0.26**	0.00	0.10
Real consumption	0.12	0.84	-0.01	0.05	0.92	0.18	0.15	0.01
Real nonresidential investment	0.58	0.31	0.00	0.61	0.26	**0.21**	0.06	0.03
Real residential investment	**0.85**	0.05	0.03	**0.96**	0.01	**0.23**	0.04	0.05
CPI inflation	-1.02	0.34	0.00	-0.55	0.57	-0.12	0.35	-0.01
Housing starts	-0.34	0.26	0.00	-0.35	0.12	0.00	0.99	0.00
Industrial production	**0.95**	0.01	0.03	**0.92**	0.10	-0.08	0.65	0.02
Unemployment rate	**0.72**	0.00	0.17	**1.02**	0.00	**-0.94**	0.00	0.21
Employment growth	**2.09**	0.00	0.28	**1.97**	0.00	**0.57**	0.00	0.32
Treasury bill 3 month rate	**0.40**	0.00	0.05	0.14	0.20	**1.38**	0.00	0.16
Treasury bond 10 year rate	-0.16	0.27	0.00	-0.15	0.30	-0.01	0.98	-0.01
AAA Corporate bond rate	0.10	0.41	0.00	**-0.26**	0.08	**1.03**	0.01	0.09

Notes: The CG model refers to estimates of (1.3), while the Revised Model refers to estimates of (1.17). All forecasts are of x_{t+2}. The variables in the top panel are from the NIPA. Data used for estimation are from 1968Q4 to 2019Q4. Standard errors are computed using the Newey-West method with automatic bandwidth selection to correct for serial correlation. Estimates that are in bold are statistically significant at the 10% level.

Table A.19. Comparing estimates of K between CG and Revised model, h=2

Variable	**Test stat**	**pval**	**CG K**	**REV K**	**Difference**
Nominal GDP	0.02	0.89	0.71	0.72	0.00
Real GDP	0.19	0.67	0.54	0.54	0.00
GDP price index inflation	2.78	0.10	0.48	0.52	**-0.04**
Real consumption	0.73	0.39	0.89	0.95	-0.06
Real nonresidential investment	0.07	0.79	0.63	0.62	0.01
Real residential investment	0.47	0.49	0.54	0.51	0.03
CPI inflation	0.32	0.57	-44.69	2.24	-46.93
Housing starts	0.00	0.96	1.51	1.53	-0.02
Industrial production	0.33	0.57	0.51	0.52	-0.01
Unemployment rate	7.84	0.01	0.58	0.50	**0.09**
Employment growth	0.68	0.41	0.32	0.34	-0.01
Treasury bill 3 month rate	6.80	0.01	0.71	0.88	**-0.17**
Treasury bond 10 year rate	2.05	0.15	1.19	1.18	0.01
AAA Corporate bond rate	8.41	0.00	0.91	1.36	**-0.45**

Notes: The test statistic is a chi-squared adjusted for heteroskedasticity. CG K and REV K refer to the K recovered from both model estimates according to the expression $K = \frac{1}{1+\hat{\beta}}$. Estimates that are in bold are statistically different at the 10% level.

Table A.20. Average Forecast Errors and Average Forecast Updates, h=3

	CG Model			Revised Model				
Variable	β	**pval**	R^2	β	**pval**	β_r	**pval**	R^2
Nominal GDP	-0.46	0.21	0.00	-0.42	0.26	**-0.22**	0.07	0.02
Real GDP	**-0.42**	0.03	0.01	-0.39	0.11	-0.11	0.30	0.01
GDP price index inflation	**1.29**	0.00	0.10	**1.24**	0.00	0.08	0.18	0.10
Real consumption	-0.42	0.31	0.00	-0.12	0.70	0.03	0.69	-0.01
Real nonresidential investment	0.99	0.12	0.01	**1.16**	0.05	-0.08	0.48	0.01
Real residential investment	-0.56	0.52	0.00	0.43	0.58	0.15	0.29	0.00
CPI inflation	-0.89	0.33	0.00	-0.24	0.75	-0.15	0.42	-0.01
Housing starts	-0.29	0.46	0.00	-0.02	0.91	**0.31**	0.00	0.04
Industrial production	-0.08	0.85	-0.01	-0.21	0.61	**-0.34**	0.03	0.03
Unemployment rate	**0.83**	0.00	0.12	**1.17**	0.00	**-1.08**	0.01	0.15
Employment growth	**2.29**	0.00	0.15	**2.28**	0.00	**0.41**	0.10	0.15
Treasury bill 3 month rate	**0.67**	0.00	0.08	**0.26**	0.08	**1.97**	0.00	0.21
Treasury bond 10 year rate	-0.06	0.42	-0.01	-0.05	0.72	**-0.55**	0.06	0.00
AAA Corporate bond rate	0.08	0.61	-0.01	**-0.30**	0.02	**1.30**	0.01	0.11

Notes: The CG model refers to estimates of (1.3), while the Revised Model refers to estimates of (1.17). All forecasts are of x_{t+3}. The variables in the top panel are from the NIPA. Data used for estimation are from 1968Q4 to 2019Q4. Standard errors are computed using the Newey-West method with automatic bandwidth selection to correct for serial correlation. Estimates that are in bold are statistically significant at the 10% level.

Table A.21. Comparing estimates of K between CG and Revised model, h=3

Variable	Test stat	pval	CG K	REV K	Difference
Nominal GDP	0.31	0.58	1.85	1.71	0.13
Real GDP	0.22	0.64	1.73	1.64	0.09
GDP price index inflation	0.52	0.47	0.44	0.45	-0.01
Real consumption	5.23	0.02	1.73	1.14	**0.60**
Real nonresidential investment	3.00	0.08	0.50	0.46	**0.04**
Real residential investment	12.40	0.00	2.28	0.70	**1.58**
CPI inflation	0.42	0.52	9.49	1.31	8.18
Housing starts	0.99	0.32	1.40	1.02	0.38
Industrial production	1.72	0.19	1.09	1.26	-0.17
Unemployment rate	5.17	0.02	0.55	0.46	**0.08**
Employment growth	0.02	0.90	0.30	0.31	0.00
Treasury bill 3 month rate	8.25	0.00	0.60	0.79	**-0.20**
Treasury bond 10 year rate	0.36	0.55	1.07	1.05	0.02
AAA Corporate bond rate	4.62	0.03	0.93	1.42	**-0.49**

Notes: The test statistic is a chi-squared adjusted for heteroskedasticity. CG K and REV K refer to the K recovered from both model estimates according to the expression $K = \frac{1}{1+\hat{\beta}}$. Estimates that are in bold are statistically different at the 10% level.

A.2 Appendix - Chapter 3

A.2.1 Proofs of propositions

Proposition 5

First I prove the proposition for an AR(2) and then demonstrate how this extends to the AR(p) case. The average forecast equation for the AR(2) state vector $\boldsymbol{z}_t$ is as follows:

$$\begin{bmatrix} \mathbb{E}_t[x_t] \\ \mathbb{E}_t[x_{t-1}] \end{bmatrix} = \begin{bmatrix} \mathbb{E}_{t-1}[x_t] \\ \mathbb{E}_{t-1}[x_{t-1}] \end{bmatrix} + \begin{bmatrix} K_1 & K_r \\ K_2 & K_{r_2} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \left(\begin{bmatrix} x_t \\ x_{t-1} \end{bmatrix} - \begin{bmatrix} \mathbb{E}_{t-1}[x_t] \\ \mathbb{E}_{t-1}[x_{t-1}] \end{bmatrix} \right)$$

where K_r, K_{r_2} are the weights placed on the current and lagged signal innovations for forecasting the current and lagged states, respectively. Averaging across individuals removes the idiosyncratic noise terms, since $\mathbb{E}_{t-1}[\psi_{it}] = \mathbb{E}_{t-1}[\nu_{it}] = 0$. The goal is derive the relationship between average forecast errors, $x_t - \mathbb{E}_t[x_t]$, and the average forecast update, $\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t]$. The average forecast of x_t is:

$$\mathbb{E}_t[x_t] = \mathbb{E}_{t-1}[x_t] + K_1(x_t - \mathbb{E}_{t-1}[x_t]) + K_r(x_{t-1} - \mathbb{E}_{t-1}[x_{t-1}])$$

Using properties of the state I can rewrite x_{t-1} as:

$$x_{t-1} = \frac{1}{\rho_1} x_t - \frac{\rho_2}{\rho_1} x_{t-2} - \frac{w_t}{\rho_1}$$

If I plug this into $\mathbb{E}_{t-1}[x_{t-1}]$ I find:

$$\mathbb{E}_{t-1}[x_{t-1}] = \frac{1}{\rho_1} E_{t-1}[x_t] - \frac{\rho_2}{\rho_1} E_{t-1}[x_{t-2}]$$

since $\mathbb{E}_{t-1}[w_t] = 0$. Using these two expressions I can rewrite the third term of the average forecast above:

$$x_{t-1} - \mathbb{E}_{t-1}[x_{t-1}] = \frac{1}{\rho_1}(x_t - \mathbb{E}_{t-1}[x_t]) - \frac{\rho_2}{\rho_1}(x_{t-2} - \mathbb{E}_{t-1}[x_{t-2}]) - \frac{w_t}{\rho_1}$$

Plugging this into the average forecast and grouping terms gives

$$\mathbb{E}_t[x_t] = \mathbb{E}_{t-1}[x_t] + \left(\frac{\rho_1 K_1 + K_r}{\rho_1}\right)(x_t - \mathbb{E}_{t-1}[x_t]) - \left(\frac{\rho_2 K_r}{\rho_1}\right)(x_{t-2} - \mathbb{E}_{t-1}[x_{t-2}]) - K_r\left(\frac{w_t}{\rho_1}\right)$$

Finally I subtract $\left(\frac{\rho_1(1-K_1)-K_r}{\rho_1}\right)\mathbb{E}_{t-1}[x_t]$ and $\left(\frac{\rho_1 K_1+K_r}{\rho_1}\right)\mathbb{E}_t[x_t]$ from both sides and rearrange to find the expression in the proposition:

$$\begin{aligned} x_t - \mathbb{E}_t[x_t] = &\left(\frac{\rho_1(1-K_1)-K_r}{\rho_1 K_1 + K_r}\right)(\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t]) \\ &+ \left(\frac{\rho_2 K_r}{\rho_1 K_1 + K_r}\right)(x_{t-2} - \mathbb{E}_{t-1}[x_{t-2}]) + \left(\frac{\rho_1 K_r}{K_1 + K_r}\right) w_t \end{aligned}$$

Note that when the fundamental follows an AR(1), this implies $\rho_2 = 0$, and the expression collapses to the same bias equation from chapter one.

To extend this case to an AR(p), first note that the average forecast equation is the same as that of an AR(2), since it does not depend on the lag length of the fundamental. Next, rewrite x_{t-1} as:

$$x_{t-1} = \rho_1^{-1}(x_t - \rho_2 x_{t-2} - \rho_3 x_{t-3} - \cdots - \rho_p x_{t-p} - w_t)$$

Substituting this into the average forecast equation and rearrange to find

$$\begin{aligned} \mathbb{E}_t[x_t] = &\left(\frac{\rho_1(1-K_1)-K_r}{\rho_1}\right)\mathbb{E}_{t-1}[x_t] + \left(\frac{\rho_1 K_1 + K_r}{\rho_1}\right) x_t \\ &- \left(\frac{\rho_2 K_r}{\rho_1}\right)(x_{t-2} - \mathbb{E}_{t-1}[x_{t-2}]) - \left(\frac{\rho_3 K_r}{\rho_1}\right)(x_{t-3} - \mathbb{E}_{t-1}[x_{t-3}]) \\ &- \cdots - \left(\frac{\rho_p K_r}{\rho_1}\right)(x_{t-p} - \mathbb{E}_{t-1}[x_{t-p}]) - \left(\frac{K_r}{\rho_1}\right) w_t \end{aligned}$$

I subtract $\left(\frac{\rho_1(1-K_1)-K_r}{\rho_1}\right)\mathbb{E}_{t-1}[x_t]$ and $\left(\frac{\rho_1 K_1+K_r}{\rho_1}\right)\mathbb{E}_t[x_t]$ from both sides, as in the AR(2) case, and rearrange to find:

$$
\begin{aligned}
x_t - \mathbb{E}_t[x_t] &= \left(\frac{\rho_1(1-K_1)-K_r}{\rho_1 K_1 + K_r}\right)(\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t]) \\
&+ \left(\frac{\rho_2 K_r}{\rho_1 K_1 + K_r}\right)(x_{t-2} - \mathbb{E}_{t-1}[x_{t-2}]) + \left(\frac{\rho_3 K_r}{\rho_1 K_1 + K_r}\right)(x_{t-3} - \mathbb{E}_{t-1}[x_{t-3}]) \\
&+ \cdots + \left(\frac{\rho_p K_r}{\rho_1 K_1 + K_r}\right)(x_{t-p} - \mathbb{E}_{t-1}[x_{t-p}]) + \left(\frac{\rho_1 K_r}{K_1 + K_r}\right) w_t
\end{aligned}
$$

When the fundamental follows an AR(p), there are $p+1$ terms alongside the average forecast update. The first p terms are backcast errors made of increasingly distant horizons from $t-2$ to $t-p$. The final term is the fundamental shock w_t.

A.2.2 ARMA representations for forecasting equations

This section describes the steps I use to find the ARMA representation of the average forecasting equations for an AR(2) when there are two signals in the measurement equation. Using the model set-up from the text, the forecast for the current state is given by:

$$
\begin{aligned}
\mathbb{E}_t[x_t] &= \mathbb{E}_{t-1}[x_t] + K_1(x_t - \mathbb{E}_{t-1}[x_t]) + K_r(x_{t-1} - \mathbb{E}_{t-1}[x_{t-1}]) \\
&= (\rho_1(1-K_1) - K_r)\mathbb{E}_{t-1}[x_{t-1}] + \rho_2(1-K_1)\mathbb{E}_{t-1}[x_{t-2}] \\
&+ (\rho_1 K_1 + K_r)x_{t-1} + \rho_2 K_1 x_{t-2} + K_1 w_t \\
&= \frac{1}{1-GL}\left(\rho_2(1-K_1)\mathbb{E}_{t-1}[x_{t-2}] + K_1 x_{t-1} + K_r(x_{t-1})\right) \qquad \text{(A.15)}
\end{aligned}
$$

The forecast for the lagged state is given by:

$$
\begin{aligned}
\mathbb{E}_t[x_{t-1}] &= \mathbb{E}_{t-1}[x_{t-1}] + K_2(x_t - \mathbb{E}_{t-1}[x_t]) + K_{r_2}(x_{t-1} - \mathbb{E}_{t-1}[x_{t-1}]) \\
&= (1 - \rho_1 K_2 - K_{r_2})\mathbb{E}_{t-1}[x_{t-1}] - \rho_2 K_2 \mathbb{E}_{t-1}[x_{t-2}] \\
&+ (\rho_1 K_2 + K_{r_2})x_{t-1} + \rho_2 K_2 x_{t-2} + K_2 w_t \\
&= \frac{1}{1+\rho_2 K_2 L}\left(H\mathbb{E}_{t-1}[x_{t-1}] + (1-H)x_{t-1} + \rho_2 K_2 x_{t-2} + K_2 w_t\right) \qquad \text{(A.16)}
\end{aligned}
$$

where $G = \rho_1(1-K_1) - K_r$ and $H = 1 - \rho_1 K_2 - K_{r_2}$. After iterating equation (A.16) backwards one period I can plug equation it into (A.15) to find:

$$
\begin{aligned}
\mathbb{E}_t[x_t] =& \frac{\rho_2(1-K_1)}{1+\phi_1 L + \phi_2 L^2}\big((1-H)x_{t-2} + \rho_2 K_2 x_{t-3} + K_2 w_{t-1}\big) \\
&+ \frac{1+\rho_2 K_2 L}{1+\phi_1 L + \phi_2 L^2}\left((\rho_1 - G)x_{t-1} + \rho_2 K_1 x_{t-2} + K_1 w_t\right) \\
=& \frac{K_1 x_t + (\rho_2 K_2 + K_r)x_{t-1} + \rho_2(K_{r_2}(1-K_1) + K_2 K_r)(x_{t-2})}{1+\phi_1 L + \phi_2 L^2} \qquad \text{(A.17)}
\end{aligned}
$$

where $\phi_1 = (\rho_2 K_2 - \rho_1(1-K_1) + K_r)$ and $\phi_2 = \rho_2(K_2 K_r - (1-K_1)(1-K_{r_2}))$. Note that when the state follows an AR(1), that implies $\rho_2 = 0$, so equation (A.17) collapses to:

$$
\mathbb{E}_t[x_t] = \frac{K + K_r L}{1 - GL} x_t
$$

which is the same equation found in chapter one. To find the forecast error, simply subtract (A.17) from x_t to find:

$$
x_t - \mathbb{E}_t[x_t] = \frac{(1-K_1)(1-\rho_1 L - \rho_2 L^2)x_t}{1+\phi_1 L + \phi_2 L^2}
$$

Next I compute the backcast equation $\mathbb{E}_{t-1}[x_t]$ by iterating equation (A.15) backwards one period and plugging it into (A.16) to find:

$$\mathbb{E}_t[x_{t-1}] = \frac{K_2 x_t + [K_1(1-K_{r_2}) + K_2(K_r - \rho_1) + K_{r_2}]x_{t-1}}{1 - \Pi_1 L - \Pi_2 L^2} + \frac{[K_r(1-\rho_1 K_2) - \rho_1 K_{r_2}(1-K_1)](x_{t-2})}{1 - \Pi_1 L - \Pi_2 L^2}$$

where $\Pi_1 = (1-K_1)(\rho_1 + \rho_2(1-\rho_1 K_2 - K_{r_2}) - K_r - \rho_2 K_2)$ and $\Pi_2 = \rho_2 K_2 G$. When the fundamental follows an AR(1), the backcast becomes:

$$\mathbb{E}_t[x_{t-1}] = \frac{K_1 + K_r L}{1 - GL} x_{t-1}$$

which is similar to the equation for $\mathbb{E}_t[x_t]$ except x_t is replaced with x_{t-1}. This is because there is no need for learning about higher-order dynamics when the fundamental has no higher-order dynamics.