INTERNET APPENDIX

for

Measuring Under- and Overreaction in Expectation Formation

Simas Kučinskas Florian S. Peters

Appendix A Existing Methods

In this section, we formally analyze existing reduced-form regressions. As in Section 2.5, we consider three reduced-form regressions that are commonly used in empirical work, repeated here for convenience:

$$e_{i,t+1} = \alpha_i + \beta_{\text{CG}} r_{it} + u_{it}$$

$$e_{i,t+1} = \alpha_i + \beta_{\text{extr.}} x_t + u_{it}$$

$$e_{i,t+1} = \alpha_i + \beta_{\text{auto.}} e_{it} + u_{it}$$

We assume that we have a sample $\{e_{i,t+1}, x_t, r_{it}\}_{t=1}^{T}$ where T denotes sample size. The proposition below summarizes how estimates of the reduced-form regressions are connected to the underlying bias coefficients.[23]

Proposition 7. *Let* $w_{j\ell} \equiv \alpha_{j\ell}^2 \sigma_j^2 / \mathbb{V}\text{ar}[x_t]$, $\tilde{w}_j \equiv \sigma_j^2 / (\sum_{k=1}^{J} \sigma_k^2)$, *and* $\breve{w}_j \equiv \alpha_{j1}^2 \sigma_j^2 / \sum_{k=1}^{J} \alpha_{k1}^2 \sigma_k^2$. *OLS estimates of the regression coefficients in Eqs.* (13), (14), *and* (12) *have the following probability limits:*

$$\text{plim}_{T\to\infty} \hat{\beta}_{\text{extr.}} = -\sum_{j=1}^{J} \sum_{\substack{\ell=1 \\ \alpha_{j,\ell-1}\neq 0}}^{+\infty} w_{j,\ell-1} \left[\frac{\text{sgn}(\alpha_{j\ell}) b_{j\ell}^{(\varepsilon)}}{\alpha_{j,\ell-1}} \right]$$

$$\text{plim}_{T\to\infty} \hat{\beta}_{\text{auto.}} = -\sum_{j=1}^{J} \tilde{w}_j \left[\text{sgn}(\alpha_{j1}) b_{j1}^{(\varepsilon)} \right] + \mathcal{O}(\{[b_{j\ell}^{(\varepsilon)}]^2, [b_{\ell}^{(\omega)}]^2\})$$

$$\text{plim}_{T\to\infty} \hat{\beta}_{\text{CG}} = -\sum_{\substack{j=1 \\ \alpha_{j1}\neq 0}}^{J} \breve{w}_j \left[\text{sgn}(\alpha_{j1}) \frac{b_{j1}^{(\varepsilon)}}{\alpha_{j1}} \right] + \mathcal{O}(\{[b_{j\ell}^{(\varepsilon)}]^2, [b_{\ell}^{(\omega)}]^2\})$$

Proof. Recall that for a generic regression of $e_{i,t+1} = \alpha_i + \beta z_{it} + u_{it}$ for some regressor z_{it}, the OLS estimator of β is given by

$$\hat{\beta} = \frac{\sum_{t=1}^{T} (e_{i,t+1} - \bar{e}_{i,t+1})(z_{it} - \bar{z}_{it})}{\sum_{t=1}^{T} (z_{it} - \bar{z}_{it})^2},$$

[23] For the Coibion-Gorodnichenko regression, bias coefficients in the proposition below refer to one-step-ahead bias coefficients.

where $\bar{e}_{i,t+1} \equiv 1/T \sum_{t=1}^{T} e_{i,t+1}$ and $\bar{z}_{it} \equiv 1/T \sum_{t=1}^{T} z_{it}$.

Extrapolation Regression. We calculate that

$$\mathbb{C}\text{ov}[e_{i,t+1}, x_t] = -\sum_{j=1}^{J} \sum_{\ell=1}^{+\infty} \text{sgn}(\alpha_{j\ell}) \alpha_{j,\ell-1} b_{j\ell}^{(\varepsilon)} \sigma_j^2$$
$$\mathbb{V}\text{ar}[x_t] = \sum_{j=1}^{J} \sum_{\ell=0}^{+\infty} \alpha_{j\ell}^2 \sigma_j^2$$

Therefore,

$$\text{plim}_{T\to\infty} \hat{\beta}_{\text{extr.}} = \frac{\mathbb{C}\text{ov}[e_{i,t+1}, x_t]}{\mathbb{V}\text{ar}[x_t]},$$

and the expression in the proposition follows after using the definition of $w_{j\ell}$.

Autocorrelation Regression. We calculate that

$$\mathbb{C}\text{ov}[e_{i,t+1}, e_{it}] = -\sum_{j=1}^{J} \text{sgn}(\alpha_{j1}) b_{j1}^{(\varepsilon)} \sigma_j^2 + \sum_{j=1}^{J} \sum_{\ell=2}^{+\infty} \text{sgn}(\alpha_{j\ell}) \, \text{sgn}(\alpha_{j,\ell-1}) b_{j\ell}^{(\varepsilon)} b_{j,\ell-1}^{(\varepsilon)} \sigma_j^2 + \sigma_\omega^2 \sum_{\ell=2}^{\infty} \text{sgn}(\gamma_\ell) \, \text{sgn}(\gamma_{\ell-1}) b_\ell^{(\omega)} b_{\ell-1}^{(\omega)}$$
$$\mathbb{V}\text{ar}[e_{it}] = \sum_{j=1}^{J} \sigma_j^2 \left(1 + \sum_{\ell=1}^{+\infty} [b_{j\ell}^{(\varepsilon)}]^2\right) + \sigma_\omega^2 \sum_{\ell=1}^{\infty} [b_\ell^{(\omega)}]^2$$

We next take a first-order Taylor expansion of plim $\hat{\beta}_{\text{auto.}}$ around $b_{j\ell}^{(\varepsilon)} = b_\ell^{(\omega)} = 0$. We first note that plim $\hat{\beta}_{\text{auto.}} = 0$ at $b_{j\ell}^{(\varepsilon)} = b_\ell^{(\omega)} = 0$ and

$$\left.\frac{\partial \, \text{plim} \, \hat{\beta}_{\text{auto.}}}{\partial b_{j\ell}^{(\varepsilon)}}\right|_{b_{j\ell}^{(\varepsilon)} = b_\ell^{(\omega)} = 0} = \frac{1}{\sum_{j=1}^{J} \sigma_j^2} \left.\frac{\partial \, \mathbb{C}\text{ov}[e_{i,t+1}, e_{it}]}{\partial b_{j\ell}^{(\varepsilon)}}\right|_{b_{j\ell}^{(\varepsilon)} = b_\ell^{(\omega)} = 0}$$
$$\left.\frac{\partial \, \text{plim} \, \hat{\beta}_{\text{auto.}}}{\partial b_{\ell}^{(\omega)}}\right|_{b_{j\ell}^{(\varepsilon)} = b_\ell^{(\omega)} = 0} = \frac{1}{\sum_{j=1}^{J} \sigma_j^2} \left.\frac{\partial \, \mathbb{C}\text{ov}[e_{i,t+1}, e_{it}]}{\partial b_{\ell}^{(\omega)}}\right|_{b_{j\ell}^{(\varepsilon)} = b_\ell^{(\omega)} = 0}$$

We now calculate that

$$\left.\frac{\partial\,\mathbb{C}\text{ov}[e_{i,t+1},e_{it}]}{\partial b_{j1}^{(\varepsilon)}}\right|_{b_{j\ell}^{(\varepsilon)}=b_{\ell}^{(\omega)}=0} = -\,\text{sgn}(\alpha_{j1})\sigma_j^2$$

$$\left.\frac{\partial\,\mathbb{C}\text{ov}[e_{i,t+1},e_{it}]}{\partial b_{j\ell}^{(\varepsilon)}}\right|_{b_{j\ell}^{(\varepsilon)}=b_{\ell}^{(\omega)}=0} = 0 \text{ for } \ell \geq 2$$

$$\left.\frac{\partial\,\mathbb{C}\text{ov}[e_{i,t+1},e_{it}]}{\partial b_{\ell}^{(\omega)}}\right|_{b_{j\ell}^{(\varepsilon)}=b_{\ell}^{(\omega)}=0} = 0$$

The result in the proposition follows after using the definition of $\tilde{w}_j$.

Coibion-Gorodnichenko Regression. We use the extension to multiple-step-ahead forecasts from the Online Appendix (Appendix B.1). We first calculate forecast revisions:

$$r_{it} = [b_{i0}^{(1)} - b_{i0}^{(2)}] + \sum_{j=1}^{J} a_{j1}^{(1)}\varepsilon_{jt} + \sum_{j=1}^{J}\sum_{\ell=2}^{+\infty}\text{sgn}(\alpha_{j\ell})[b_{j\ell}^{(\varepsilon,1)} - b_{j\ell}^{(\varepsilon,2)}]\varepsilon_{j,t+1-\ell} + [\nu_{it}^{(1)} - \nu_{i,t-1}^{(2)}].$$

We next note that

$$\nu_{it}^{(1)} - \nu_{i,t-1}^{(2)} = \text{sgn}(\gamma_1^{(1)})b_1^{(\omega,1)}\omega_{it} + \sum_{\ell=2}^{+\infty}[\text{sgn}(\gamma_{\ell}^{(1)})b_{\ell}^{(\omega,1)} - \text{sgn}(\gamma_{\ell-1}^{(2)})b_{\ell-1}^{(\omega,2)}]\omega_{i,t-\ell}.$$

Therefore, we deduce that

$$\begin{aligned}\mathbb{C}\text{ov}[e_{i,t+1},r_{it}] = &-\sum_{j=1}^{J}\text{sgn}(\alpha_{j1})b_{j1}^{(\varepsilon,1)}[\alpha_{j1}+\text{sgn}(\alpha_{j1})b_{j1}^{(\varepsilon,1)}]\sigma_j^2 - \sum_{j=1}^{J}\sum_{\ell=2}^{+\infty}b_{j\ell}^{(\varepsilon,1)}[b_{j\ell}^{(\varepsilon,1)} - b_{j\ell}^{(\varepsilon,2)}]\sigma_j^2 \\ &- [b_1^{(\omega,1)}]^2\sigma_\omega^2 - \sigma_\omega^2\sum_{\ell=2}^{+\infty}\text{sgn}(\gamma_{\ell}^{(1)})b_{\ell}^{(\omega,1)}[\text{sgn}(\gamma_{\ell}^{(1)})b_{\ell}^{(\omega,1)} - \text{sgn}(\gamma_{\ell-1}^{(2)})b_{\ell-1}^{(\omega,2)}]\end{aligned}$$

$$\begin{aligned}\mathbb{V}\text{ar}[r_{it}] = &\sum_{j=1}^{J}[\alpha_{j1}+\text{sgn}(\alpha_{j1})b_{j1}^{(\varepsilon,1)}]^2\sigma_j^2 + \sum_{j=1}^{J}\sum_{\ell=2}^{+\infty}[b_{j\ell}^{(\varepsilon,1)} - b_{j\ell}^{(\varepsilon,2)}]^2\sigma_j^2 \\ &+ [b_1^{(\omega,1)}]^2\sigma_\omega^2 + \sigma_\omega^2\sum_{\ell=2}^{+\infty}[\text{sgn}(\gamma_{\ell}^{(1)})b_{\ell}^{(\omega,1)} - \text{sgn}(\gamma_{\ell-1}^{(2)})b_{\ell-1}^{(\omega,2)}]^2\end{aligned}$$

We again consider the Taylor expansion around $b_{j\ell}^{(\varepsilon,h)} = b_{\ell}^{(\omega,h)} = 0$. We first note that plim $\hat{\beta}_{\text{CG}} =$

0 at $b_{j\ell}^{(\varepsilon,h)} = b_{\ell}^{(\omega,h)} = 0$. Next, we have that

$$\left.\frac{\partial \operatorname{plim} \hat{\beta}_{\text{CG}}}{\partial b_{j\ell}^{(\varepsilon,h)}}\right|_{b_{j\ell}^{(\varepsilon,h)}=b_{\ell}^{(\omega,h)}=0} = \frac{1}{\sum_{j=1}^{J} \alpha_{j1}^2 \sigma_j^2} \left.\frac{\partial \mathbb{C}\text{ov}[e_{i,t+1}, r_{it}]}{\partial b_{j\ell}^{(\varepsilon,h)}}\right|_{b_{j\ell}^{(\varepsilon,h)}=b_{\ell}^{(\omega,h)}=0}$$
$$\left.\frac{\partial \operatorname{plim} \hat{\beta}_{\text{CG}}}{\partial b_{\ell}^{(\omega,h)}}\right|_{b_{j\ell}^{(\varepsilon,h)}=b_{\ell}^{(\omega,h)}=0} = \frac{1}{\sum_{j=1}^{J} \alpha_{j1}^2 \sigma_j^2} \left.\frac{\partial \mathbb{C}\text{ov}[e_{i,t+1}, r_{it}]}{\partial b_{\ell}^{(\omega,h)}}\right|_{b_{j\ell}^{(\varepsilon,h)}=b_{\ell}^{(\omega,h)}=0}$$

We now calculate that

$$\left.\frac{\partial \operatorname{plim} \hat{\beta}_{\text{CG}}}{\partial b_{j1}^{(\varepsilon,1)}}\right|_{b_{j\ell}^{(\varepsilon,h)}=b_{\ell}^{(\omega,h)}=0} = -\operatorname{sgn}(\alpha_{j1})\alpha_{j1}\sigma_j^2$$
$$\left.\frac{\partial \operatorname{plim} \hat{\beta}_{\text{CG}}}{\partial b_{j\ell}^{(\varepsilon,h)}}\right|_{b_{j\ell}^{(\varepsilon,h)}=b_{\ell}^{(\omega,h)}=0} = 0 \text{ for } \ell \geq 2$$
$$\left.\frac{\partial \operatorname{plim} \hat{\beta}_{\text{CG}}}{\partial b_{\ell}^{(\omega,h)}}\right|_{b_{j\ell}^{(\varepsilon,h)}=b_{\ell}^{(\omega,h)}=0} = 0$$

The result follows after substituting the expression for $\breve{w}_j$. □

For autocorrelation and Coibion-Gorodnichenko regressions, we use a first-order approximation to obtain readily-interpretable expressions. This approximation is accurate as long as biases are "small" (i.e., $b_{j\ell}^{(\varepsilon)} \approx 0$ and $b_{\ell}^{(\omega)} \approx 0$). Exact closed-form expressions for the autocorrelation and Coibion-Gorodnichenko regressions can be found in the proof above.[24]

Intuition. Proposition 7 demonstrates that commonly-used reduced-form regressions essentially estimate a weighted average of the underlying bias coefficients. The underlying weights differ across the reduced-form regressions. For the extrapolation regression, the weights are given by the fraction of the variance of x_t due to shock j at lag ℓ. For the autocorrelation regression, bias coefficients are weighted by the share of the variance of $\sum_{k=1}^{J} \varepsilon_{kt}$ (i.e., the combined time-t innovation) due to shock j. Finally, for the Coibion-Gorodnichenko regres-

[24] Proposition 7 refers to regressions estimated forecaster-level data. Up to a first-order approximation, the results are identical for consensus-level regressions. To obtain corresponding results for the exact closed-form expressions at the consensus level, simply set all of the idiosyncratic bias coefficients to zero.

sions, the weights are given by the fraction of the variance of $\mathbb{E}_t[x_{t+1}] - \mathbb{E}_{t-1}[x_t]$ (i.e., the full-information rational expectations revision) due to shock j. Hence, reduced-form approaches may yield conflicting results simply because they weight under- and overreaction to various shocks differently.

The autocorrelation regression estimates a weighted average of the underlying bias coefficients. In contrast, the extrapolation and Coibion-Gorodnichenko regressions implicitly normalize the bias coefficients by the corresponding IRF of the target variable. Hence, these two regressions, in effect, estimate a weighted average of the underlying *percentage* bias coefficients. Such normalization by the IRF of the target variable can help interpretability. However, it is important to recognize that these implicit normalizations differ across methods.

In empirical work, positive (negative) estimates from these reduced-form regressions are typically interpreted as evidence for underreaction (overreaction). Unfortunately, for extrapolation and autocorrelation regressions, the estimated coefficients can be negative even when expectations exhibit underreaction (or positive when expectations exhibit overreaction). The reason, as Proposition 7 demonstrates, is that these regressions estimate bias coefficients multiplied by the sign of the true IRF. As a result, these regressions conflate bias coefficients with properties of the underlying data-generating process.

To understand this, assume that expectations underreact at all lags, so that all bias coefficients are negative. In addition, suppose that the process for x_t exhibits "reversals" in the sense that a positive shock to x_t today reduces its value in the next period. A positive shock to x_t implies a high value of the right-hand side variable in both extrapolation and autocorrelation regressions. Since x_t exhibits reversals, the fully-rational forecast for time $(t+1)$ is lower. However, since expectations underreact, the adjustment is not sufficient, resulting in a negative forecast error at time $(t+1)$, i.e., a low value of the left-hand side variable. Therefore, although expectations underreact, the extrapolation and autocorrelation coefficients are negative. As this example shows, extrapolation and autocorrelation regressions can misleadingly suggest

overreaction in the presence of reversals.[25]

Appendix B Extensions

B.1 Multiple-Step-Ahead Forecasts

Suppose that the agent makes h-step ahead forecasts ($h \geq 1$). We generalize the expectation-formation process in Eq. (2) to

$$\mathbb{F}_{it}[x_{t+h}] = b_{i0}^{(h)} + \sum_{j=1}^{J} \sum_{\ell=h}^{+\infty} a_{j\ell}^{(h)} \varepsilon_{j,t+h-\ell} + \nu_{it}^{(h)},$$

where $\nu_{it}^{(h)} = \sum_{\ell=1}^{+\infty} \gamma_{\ell}^{(h)} \omega_{i,t+1-\ell}$ is the idiosyncratic component. We allow the coefficients $b_{i0}^{(h)}$, $a_{j\ell}^{(h)}$, and $\gamma_{\ell}^{(h)}$ to depend on the forecast horizon.

As in Section 2.1, we define idiosyncratic bias coefficients as $b_{\ell}^{(\omega,h)} \equiv |\gamma_{\ell}^{(h)}|$. We generalize the definition of aggregate bias coefficients to $b_{j\ell}^{(\varepsilon,h)} \equiv \text{sgn}(\alpha_{j,\ell+h-1})(a_{j,\ell+h-1}^{(h)} - \alpha_{j,\ell+h-1})$. With this definition, $b_{j\ell}^{(\varepsilon,h)}$ measures how expectations at time t under/overreact to $\varepsilon_{j,t+1-\ell}$. For example, $b_{j1}^{(\varepsilon,h)}$ measures how expectations respond to the most recent realization of ε_{jt}.

The same calculations as in Section 2.1 show that h-step-ahead forecast errors, $e_{i,t+h}^{(h)} \equiv x_{t+h} - \mathbb{F}_{it}[x_{t+h}]$, follow

$$\begin{aligned} e_{i,t+h}^{(h)} = -b_{i0}^{(h)} &- \sum_{j=1}^{J} \sum_{\ell=1}^{+\infty} \text{sgn}(\alpha_{j,\ell+h-1}) b_{j\ell}^{(\varepsilon,h)} \varepsilon_{j,t+1-\ell} \\ &- \sum_{\ell=1}^{\infty} \text{sgn}(\gamma_{\ell}^{(h)}) b_{\ell}^{(\omega,h)} \omega_{i,t-\ell} + \sum_{j=1}^{J} \sum_{\ell=0}^{h-1} \alpha_{j\ell} \varepsilon_{j,t+h-\ell}. \end{aligned}$$

As before, bias coefficients are identified (up to sign) from the IRF of forecast errors. The only difference from the case studied in Section 2.1 is that the first $(h-1)$ impulse responses of forecast errors are mechanical. As is well known, even if subjective expectations react to news

[25] The Coibion-Gorodnichenko regression is not vulnerable to this issue, as can be seen by investigating the expression given in the proposition. Intuitively, the reason is that the sign of the forecast revision implicitly incorporates the sign of the true IRF.

in an unbiased way, forecast errors are mechanically autocorrelated—up to lag $(h-1)$—if the underlying process is autocorrelated at these lags (see, e.g. Diebold and Lopez, 1996). Hence, when estimating bias coefficients for multiple-step-ahead forecasts empirically, care must be taken to discard these mechanical impulse responses.

B.2 Heterogeneity in Expectation Formation

Existing research has documented that expectations are heterogeneous across individuals (e.g., Manski, 2004, Section 5). Our method can easily be applied to subsamples of the population. For example, we may estimate the bias coefficients for young and old forecasters separately. When the method is applied to average (or consensus) forecasts, it recovers an average of individual-specific bias coefficients, as we now show.

We generalize the expectation-formation process in Eq. (2) to

$$\mathbb{F}_{it}[x_{t+1}] = b_{i0} + \sum_{j=1}^{J}\sum_{\ell=1}^{+\infty} a_{ij\ell}\varepsilon_{j,t+1-\ell} + \nu_{it} \text{ for } i = 1, 2, \ldots, I,$$

where $\nu_{it} = \sum_{\ell=1}^{\infty} \gamma_{i\ell}\omega_{i,t+1-\ell}$ Specifically, the response to aggregate and idiosyncratic shocks is allowed to vary across forecasters. Denote the consensus (average) forecast by $\bar{\mathbb{F}}_t[x_{t+1}] \equiv \frac{1}{I}\sum_{i=1}^{I}\mathbb{F}_{it}[x_{t+1}]$. Then, the consensus forecast is given by

$$\bar{\mathbb{F}}_t[x_{t+1}] = \bar{b}_0 + \sum_{j=1}^{J}\sum_{\ell=1}^{+\infty} \bar{a}_{j\ell}\varepsilon_{i,t+1-\ell} + \bar{\nu}_t,$$

where $\bar{b}_0 \equiv \frac{1}{I}\sum_{i=1}^{I} b_{i0}$, $\bar{a}_{j\ell} \equiv \frac{1}{I}\sum_{i=1}^{I} a_{ij\ell}$, and $\bar{\nu}_t \equiv \frac{1}{I}\sum_{i=1}^{I}\nu_{it}$. Aggregate bias coefficients of the consensus forecast, $\bar{b}^{(\varepsilon)}_{j\ell}$, are equal to

$$\bar{b}^{(\varepsilon)}_{j\ell} = \text{sgn}(\alpha_{j\ell})(\bar{a}_{j\ell} - \alpha_{j\ell}) = \text{sgn}(\alpha_{j\ell})\left[\frac{1}{I}\sum_{i=1}^{I}\{a_{ij\ell} - \alpha_{j\ell}\}\right] = \frac{1}{I}\sum_{i=1}^{I} b^{(\varepsilon)}_{ij\ell},$$

where $b^{(\varepsilon)}_{ij\ell} \equiv \text{sgn}(\alpha_{j\ell})[a_{ij\ell} - \alpha_{j\ell}]$ are individual-specific aggregate bias coefficients. Therefore, with heterogeneity, aggregate bias coefficients of the consensus forecast equal the average aggregate bias coefficients of the individual forecasts. Hence, the IRF of forecast errors of the

consensus forecast identifies (up to sign) an average of individual-specific bias coefficients.

Given our assumptions, $\bar{\nu}_t$ converges to zero in probability as $I \to +\infty$ by the Law of Large Numbers. Hence, for large cross-sections ($I \to +\infty$), the consensus forecast does not exhibit any overreaction to idiosyncratic shocks.

B.3 State and Time Dependence

It is straightforward to generalize our framework to allow for state and time dependence in both (i) the data-generating process for the target variable; and (ii) the expectations-formation process.

Instead of being time invariant, suppose now that the true reactions $\alpha_{j\ell}$ are given by a function of a vector of state variables $\mathbf{s}_t$, so that $\alpha_{j\ell} = \alpha_{j\ell}(\mathbf{s}_t)$, with $\alpha_{j\ell}(\cdot)$ a deterministic function mapping the state variables to the true reaction of the target variable.[26] Various types of state dependence can be achieved by choosing specific state variables $\mathbf{s}_t$ and functional forms for $\alpha_{j\ell}(\cdot)$. Similarly, let the perceived reaction now be $a_{j\ell} = a_{j\ell}(\mathbf{s}_t)$. In that case, aggregate bias coefficients are also state dependent and given by

$$b_{j\ell}^{(\varepsilon)}(\mathbf{s}_t) = \text{sgn}[\alpha_{j\ell}(\mathbf{s}_t)][a_{j\ell}(\mathbf{s}_t) - \alpha_{j\ell}(\mathbf{s}_t)].$$

Hence, state dependence does not affect the basic theory in any substantial way. The only difference is that estimating state-dependent bias coefficients in practice now requires controlling for the underlying state variables appropriately.

Example. For an explicit example of state-dependent bias coefficients, consider the setting studied by Patton and Timmermann (2007, pp. 898–899). A forecaster chooses her forecast to

[26] If $\alpha_{j\ell}(\cdot)$ is not a deterministic function, then bias coefficients become random variables (even conditional on $\mathbf{s}_t$). In such cases, it may be desirable to estimate the whole distribution of bias coefficients, instead of just their first moment. See Koop, Pesaran and Potter (1996) for a discussion of the issues.

minimize the linex loss function

$$L(e_t;\kappa) = \frac{1}{\kappa^2}[\exp(\kappa e_t) - \kappa e_t - 1],$$

where κ is a parameter controlling asymmetry in losses. If $\kappa \to 0$, the loss function is proportional to e^2, and hence symmetric around zero. If $\kappa > 0$, the agent dislikes positive forecast errors (i.e., x_t is higher than predicted) more than negative forecast errors, and vice versa when $\kappa < 0$.

The true process for x_t is given by

$$\begin{aligned} x_t &= \rho x_{t-1} + \varepsilon_t \\ \varepsilon_t &= \sigma_t v_t, v_t \overset{\text{iid}}{\sim} \mathcal{N}(0,1) \\ \sigma_t^2 &= \Omega + \beta_0 \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 \end{aligned}$$

with $\rho \in (0,1)$, $\Omega > 0$, and $\beta_0, \beta_1 \geq 0$. In contrast to the setting studied in Section 2.4, aggregate shocks ε_t now follow a GARCH(1, 1) process and are therefore conditionally heteroskedastic. To ensure that the $\{\varepsilon_t\}$ process is stationary, assume that $\beta_0 + \beta_1 < 1$.

Some calculus shows that the optimal forecast is given by

$$\mathbb{F}_t[x_{t+1}] = \rho x_t + \frac{\kappa}{2}\sigma_{t+1}^2,$$

where for simplicity we omit the i subscript. Iterating the law of motion for σ_t^2, we have

$$\sigma_t^2 = \frac{\Omega}{1-\beta_1} + \beta_0 \sum_{\ell=0}^{+\infty} \beta_1^\ell \varepsilon_{t-1-\ell}^2,$$

and therefore

$$\begin{aligned} \mathbb{F}_t[x_{t+1}] &= \rho x_t + \frac{\kappa}{2}\left(\frac{\Omega}{1-\beta_1} + \beta_0 \sum_{\ell=0}^{+\infty} \beta_1^\ell \varepsilon_{t-\ell}^2\right) \\ &= \frac{\kappa\Omega}{2(1-\beta_1)} + \sum_{\ell=0}^{+\infty}\left(\rho^{\ell+1} + \frac{\kappa}{2}\beta_0\beta_1^\ell \varepsilon_{t-\ell}\right)\varepsilon_{t-\ell}. \end{aligned}$$

We therefore have

$$a_\ell(\mathbf{s}_t) = \rho^\ell + \frac{\kappa}{2}\beta_0\beta_1^{\ell-1}\varepsilon_{t-1-\ell}$$
$$b_\ell(\mathbf{s}_t) = \frac{\kappa}{2}\beta_0\beta_1^{\ell-1}\varepsilon_{t-1-\ell}$$

with the state variable $\mathbf{s}_t$ given by all past shocks, i.e., $\mathbf{s}_t = (\varepsilon_t, \varepsilon_{t-1}, \dots)^\top$. Depending on the realized values of the shocks, the model can generate both under- and overreaction. This example also illustrates how the first-order autocorrelation can be misleading as a measure of underreaction. Patton and Timmermann (2007, Eq. (23), p. 899) show that for this model, the first-order autocorrelation of forecast errors is always positive. Nevertheless, the model can in fact generate both under- and overreaction.

B.4 Measuring Deviations from Bayesian Updating

Our framework in Section 2 is designed to measure deviations from full-information rational expectations. As a result, non-zero bias coefficients should *not* be taken to imply that expectations are irrational. Non-zero bias coefficients may be due to deviations from rationality, but they may also stem from imperfect information, strategic forecaster behavior, or other frictions. For example, in the model studied in Section 2.4, expectations exhibit underreaction to the aggregate shock even with correctly perceived parameters and no diagnosticity as long as there is an information friction.

In general, measuring deviations from Bayesian updating is challenging, and other methods may be better suited for this task.[27] In particular empirical contexts, however, our method may be fruitfully applied to measure deviations from Bayesian updating.

Suppose that Bayesian forecasts can be written as

$$\mathbb{F}_t^*[x_{t+1}] = \sum_{j=1}^{J}\sum_{\ell=1}^{+\infty} a_{j\ell}^*\varepsilon_{j,t+1-\ell} + \nu_{it}^*$$

for some parameters $a_{j\ell}^*$ and γ_ℓ^*, where $\nu_{it}^* = \sum_{\ell=1}^{+\infty}\gamma_\ell^*\omega_{i,t+1-\ell}$. Here, we do not explicitly specify

[27] See, for example, Augenblick and Rabin (2021) and Augenblick and Lazarus (2018).

the underlying information structure. If the Bayesian forecaster observes the full history of ε_{jt}, $a^*_{j\ell} = \alpha_{j\ell}$ and $\gamma^*_\ell = 0$; however, that is not necessarily the case if the Bayesian forecaster has imperfect information and only observes noisy signals of the underlying aggregate shocks.

We define under- and overreaction *relative to Bayesian updating* as follows.

Definition 4. *Suppose that* $\text{sgn}(\alpha_{j\ell}) = \text{sgn}(a_{j\ell})$. *Then, expectations are said to* overreact *to aggregate shock* j *at lag* ℓ relative to Bayesian updating *if* $|a_{j\ell}| > |a^*_{j\ell}|$ *and* underreact *if* $|a_{j\ell}| < |a^*_{j\ell}|$.

Definition 5. *Suppose that* $\text{sgn}(\gamma_\ell) = \text{sgn}(\gamma^*_\ell)$. *Expectations are said to overreact to idiosyncratic shocks at lag* ℓ relative to Bayesian updating *if* $|\gamma_\ell| > |\gamma^*_\ell|$ *and underreact if* $|\gamma_\ell| < |\gamma^*_\ell|$.

We generalize the definition of bias coefficients *relative to Bayesian updating* as follows.

Definition 6. Aggregate bias coefficients relative to Bayesian updating *are defined as* $b^{(\varepsilon)}_{j\ell} \equiv \text{sgn}(\alpha_{j\ell})(a_{j\ell} - a^*_{j\ell})$, *while* idiosyncratic bias coefficients relative to Bayesian updating *are defined as* $b^{(\omega)}_{\ell} \equiv \text{sgn}(\gamma^*_\ell)(\gamma_\ell - \gamma^*_\ell)$.

We now note that the deviation of expectations from Bayesian forecasts can be written as

$$\mathbb{F}^*_t[x_{t+1}] - \mathbb{F}_{it}[x_{t+1}] = -b_{i0} - \sum_{j=1}^{J}\sum_{\ell=1}^{+\infty} \text{sgn}(\alpha_{j\ell}) b^{(\varepsilon)}_{j\ell} \varepsilon_{j,t+1-\ell} - \sum_{\ell=1}^{+\infty} \text{sgn}(\gamma^*_\ell) b^{(\omega)}_{\ell} \omega_{i,t-\ell}.$$

Hence, similarly to Proposition 1, bias coefficients relative to Bayesian updating are identified (up to sign) by the IRF of deviations of expectations from Bayesian forecasts.

To estimate these bias coefficients empirically, it is necessary to first obtain (or approximate) Bayesian forecasts. That can be readily done in experimental settings in which the Bayesian benchmark is known.[28] With observational data, the Bayesian benchmark is typically unknown. In data-rich environments, it may be possible to approximate Bayesian forecasts using machine learning (de Silva and Thesmar, 2021). In some cases, theory may give precise predictions for Bayesian forecasts. For instance, classic asset pricing theory predicts that stock prices follow martingales at short horizons (see, e.g. Cochrane, 2005, p. 22). Therefore, $x_t = \varepsilon_t$ for a martingale difference sequence ε_t is a reasonable first approximation for stock returns, implying

[28] For an example application using our method, see He and Kucinskas (2020).

$\mathbb{F}_t^*[x_{t+1}] = 0$. In this special case, valid estimates of bias coefficients relative to Bayesian updating can be obtained by simply regressing subjective expectations on past stock returns, as is done in a large literature studying expectations of stock market returns (see, e.g., Graham and Harvey, 2001; Vissing-Jorgensen, 2004; Ben-David, Graham and Harvey, 2013; Dominitz and Manski, 2011; Greenwood and Shleifer, 2014).

B.5 Measurement Error

Expectations are typically measured with some error. Measurement error can arise for many reasons, including survey design imperfections and limited attention of survey participants.[29] The magnitude of measurement error is often substantial, especially for individual-level data. For example, Giglio, Maggiori, Stroebel and Utkus (2021) elicit expected stock market returns in a sample of individual investors using two independent measurement methods. They find a correlation between the two measures of around 0.50, suggesting substantial measurement error.

The effect of measurement error depends on the type of bias coefficients being estimated. For aggregate bias coefficients, measurement error is not problematic. From the results in Section 3, estimation using an external instrument yields consistent estimates of aggregate bias coefficients as long as measurement error in expectations is uncorrelated with the instrument for the structural shock. (Of course, in finite samples measurement error reduces the precision of the estimates.) This robustness to measurement error is a key advantage of the IV procedure. However, measurement error leads to an attenuation bias when estimating composite and idiosyncratic bias coefficients, as we now discuss.

[29] A closely related possibility is that subjective expectations are inherently noisy, such as in the model of da Silveira and Woodford (2019). The empirical predictions of measurement error and inherently noisy expectations are often very similar. Noise in expectations may also indicate that expectations are not very well formed and only weakly related to actual behavior; see Drerup, Enke and von Gaudecker (2017) for more discussion and empirical evidence on this latter point.

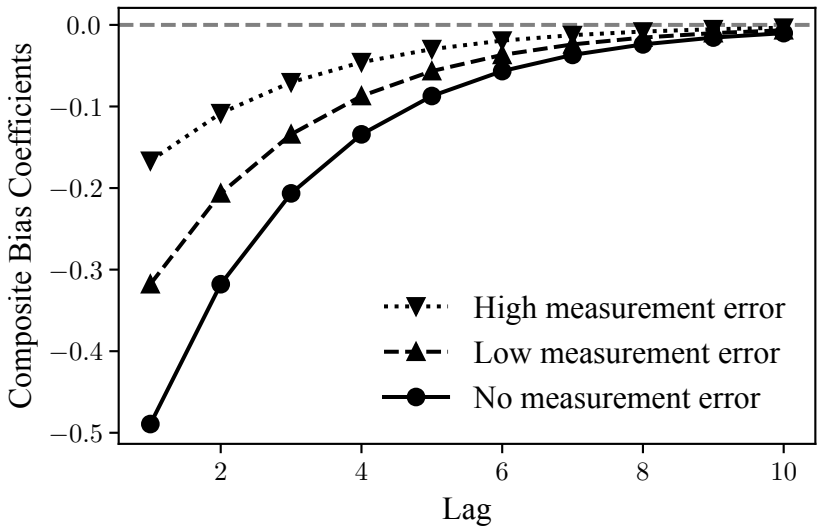

Figure A.1: Composite bias coefficients under measurement error. Same parameter values as for Figure 1. *No measurement error*: Noisy-diagnostic model. *Low measurement error*: Noisy-diagnostic model with $\sigma_m = 1.0$. *High measurement error*: Noisy-diagnostic model with $\sigma_m = 2.0$.

B.5.1 Composite Bias Coefficients

Estimating composite bias coefficients boils down to estimating the univariate IRF of forecast errors. Classical measurement error tends to mask any existing predictability of the forecast errors, biasing the estimated bias coefficients towards zero. As a result, the empirically estimated composite bias coefficients are likely to represent a lower bound on the true composite bias coefficients. To the extent that measurement error is idiosyncratic, measurement error is less problematic for consensus forecasts.

Numerical Illustration. We illustrate the effect of measurement error on composite bias coefficients in the model studied in Section 2.4. Suppose that true (but unobserved) expectations are generated as in Eq. (11). However, instead of observing the true expectations $\mathbb{F}_{it}[x_{t+1}]$, we only observe

$$\tilde{\mathbb{F}}_{it}[x_{t+1}] = \mathbb{F}_{it}[x_{t+1}] + m_{it},$$

where m_{it} is normally-distributed measurement error with mean zero and variance σ_m^2 (uncorrelated with ε_t and ω_{it} at all leads and lags). Define observed forecast errors as $\tilde{e}_{i,t+1} =$

$x_{t+1} - \tilde{\mathbb{F}}_{it}[x_{t+1}]$. Observed forecast errors are related to the true forecast errors, $e_{i,t+1}$, by

$$e_{i,t+1} = \tilde{e}_{i,t+1} + m_{it}.$$

From Proposition 4, $e_{i,t+1}$ follows an ARMA(2,2) process. Measurement error adds an additional independent MA(2) term to the observed forecast errors. Hence, observed forecast errors also follow an ARMA(2,2) process but with different parameters for the MA(2) component. These parameters can be calculated numerically given parameter values for the data-generating process.

Figure A.1 plots composite bias coefficients at different levels of measurement error. The baseline with no measurement error shows composite bias coefficients for the noisy-diagnostic model of Bordalo, Gennaioli, Ma and Shleifer (2020), as in Section 2.4 (i.e., identical to composite bias coefficients in the top-right panel of Figure 1). As seen in the graph, measurement error can lead to a substantial attenuation bias. With $\sigma_m = 1.0$ (i.e., the standard deviation of measurement error is three times smaller than σ_ω, the level of noise in the idiosyncratic signals), composite bias coefficients are reduced by around 40% relative to the case without measurement error. With $\sigma_m = 2.0$, attenuation bias is even stronger, more than halving the composite bias coefficients.

We note that although measurement error leads to an attenuation bias, this bias is conservative in the sense that measurement error makes it more difficult to find under- or overreaction in expectations. Hence, if one finds non-zero composite bias coefficients in an empirical setting in which measurement error is likely, it is probable that the true bias coefficients are even larger.

B.5.2 Idiosyncratic Bias Coefficients

Measurement error also leads to an attenuation bias when estimating idiosyncratic bias coefficients using the instrumental-variables procedure of Section 3.4. Recall that when measurement error is not present, idiosyncratic bias coefficients can be estimated consistently with the fol-

lowing IV regression:

$$e_{i,t+\ell} = \alpha_i + \beta\, \mathbb{F}_{it}[x_{t+1}] + \sum_{k=1}^{K} \pi_k d_{i,t-k} + u_{i,t+\ell}. \tag{A.1}$$

Here, for each horizon $\ell = 1, ..., L$, we instrument $\mathbb{F}_{it}[x_{t+1}]$ with the deviation from consensus, d_{it}, while lagged deviations, $d_{i,t-k}$, serve as their own instruments. The idiosyncratic bias coefficients are estimated by $\hat{b}_\ell^{(\omega)} = |\hat{\beta}|$.

Similarly to the previous section, however, suppose that instead of observing the true expectations $\mathbb{F}_{it}[x_{t+1}]$, we can only observe

$$\tilde{\mathbb{F}}_{it}[x_{t+1}] = \mathbb{F}_{it}[x_{t+1}] + m_{it},$$

where m_{it} is i.i.d. measurement error with zero mean and finite variance (independent across both time and forecasters). In that case, observed deviations from consensus, $\tilde{d}_{it}$, equal $\tilde{d}_{it} = d_{it} + m_{it}$, where d_{it} denotes the true (but unknown) deviation from consensus. Hence, measurement error introduces a mechanical positive correlation between the observed forecasts, $\tilde{\mathbb{F}}_{it}[x_{t+1}]$, and deviations from consensus, $\tilde{d}_{it}$. At the same time, since measurement error is i.i.d., m_{it} is mechanically negatively correlated with the observed forecast error at time $t+1$, but uncorrelated with forecast errors at all other leads and lags. Hence, roughly speaking, measurement error makes the first-stage regression associated with Eq. (A.1) stronger but leaves the corresponding reduced-form regression unchanged. The combined effect is that IV estimates are attenuated towards zero.

We can formalize this intuition using ideas from in Section 3.4. In particular, consider the following regression:

$$\tilde{e}_{i,t+\ell} = \alpha_i + \beta\tilde{\mathbb{F}}_{it}[x_{t+1}] + \sum_{k=1}^{K} \pi_k \tilde{d}_{i,t-k} + u_{i,t+\ell}, \tag{A.2}$$

where variables with tildes denote values observed with measurement error. As before, we instrument $\tilde{\mathbb{F}}_{it}[x_{t+1}]$ with the observed deviation from consensus, $\tilde{d}_{it}$, and lagged observed deviations, $\tilde{d}_{i,t-k}$, serve as their own instruments. We have the following result.

Proposition 8. *Consider the setup in Proposition 6. Suppose, however, that expectations are observed with i.i.d. measurement error that has zero mean and finite variance. Then, the probability limit of the IV estimate of β in Eq. (A.2) is given by*

$$\text{plim}_{T\to\infty}\,\hat{\beta} = \begin{cases} -1 & \text{if } \ell = 1 \\ -\gamma_1^{-1}\,\text{sgn}(\gamma_\ell) b_\ell^{(\omega)} \cdot K & \text{if } \ell \geq 2 \end{cases}$$

where $K \in (0,1)$ is an attenuation factor.

Proof. We closely follow the proof of Proposition 6 and notation therein. First, note that $d_{it} = (b_{i0} - \bar{b}_0) + \nu_{it} + m_{it}$, and therefore $\tilde{d}_{it}^{\perp} = \gamma_1 \omega_{it}^{\perp} + m_{it}^{\perp}$. Since $\tilde{e}_{i,t+\ell}^{\perp} = e_{i,t+\ell}^{\perp} - m_{i,t+\ell-1}^{\perp}$, we calculate that

$$\mathbb{E}[e_{i,t+\ell}^{\perp} d_{it}^{\perp}] = \begin{cases} -\gamma_1^2\,\mathbb{E}[(\omega_{it}^{\perp})^2] - \mathbb{E}[(m_{it}^{\perp})^2] & \text{if } \ell = 1 \\ -\gamma_1\gamma_\ell\,\mathbb{E}[(\omega_{it}^{\perp})^2] & \text{if } \ell \geq 2 \end{cases}$$

Letting $\tilde{f}_{it} \equiv \tilde{\mathbb{F}}_{it}[x_{t+1}]$, we find that $\mathbb{E}[\tilde{f}_{it}^{\perp}\tilde{d}_{it}^{\perp}] = \gamma_1^2\,\mathbb{E}[(\omega_{it}^{\perp})^2] + \mathbb{E}[(m_{it}^{\perp})^2]$. Hence, we obtain

$$\hat{\beta} \xrightarrow{p} \frac{\mathbb{E}[\tilde{e}_{i,t+\ell}^{\perp}\tilde{d}_{it}^{\perp}]}{\mathbb{E}[\tilde{f}_{it}^{\perp}\tilde{d}_{it}^{\perp}]} = \begin{cases} -1 & \text{if } \ell = 1 \\ -\gamma_1^{-1}\gamma_\ell \cdot K & \text{if } \ell \geq 2 \end{cases}$$

where $K \equiv \gamma_1^2\,\mathbb{E}[(\omega_{it}^{\perp})^2]/\{\gamma_1^2\,\mathbb{E}[(\omega_{it}^{\perp})^2] + \mathbb{E}[(m_{it}^{\perp})^2]\}$, and the result follows. □

The key takeaway from Proposition 8 is that in the presence of measurement error, IV estimates of idiosyncratic bias coefficients are attenuated towards zero.[30] Intuitively, with the IV procedure of Section 3.4, idiosyncratic bias coefficients are identified by the IRF of deviations from consensus. If deviations from consensus are serially correlated, the IV procedure

[30] The probability limit of the estimate at the first lag is (-1), just as in the case when no measurement error is present. The reason is that, even without measurement error, idiosyncratic bias coefficients are only estimated up to a scaling factor of γ_1^{-1}, effectively normalizing the first idiosyncratic bias coefficient to unity.

interprets this feature of the data as persistent overreaction to idiosyncratic shocks. However, classical measurement error reduces the persistence of deviations from consensus, thereby attenuating estimated idiosyncratic bias coefficients to zero. Thankfully, this attenuation bias is again conservative, with measurement error making it *less* likely that we find overreaction to idiosyncratic shocks.

Appendix C Data Appendix

C.1 Data Construction

We download data for individual inflation forecasts from the website of the Federal Reserve Bank of Philadelphia (link). The downloaded file contains forecasts of GDP deflator for the past quarter (*PGDP1*), current quarter (*PGDP2*), and the next four quarters (*PGDP3* up to *PGDP6*), see Federal Reserve Bank of Philadelphia (2017, pp. 20–22).

To construct one-quarter-ahead consensus inflation forecasts, we first calculate the median forecast of *PGDP2* and *PGDP3* in each quarter. Then, we calculate annualized quarter-on-quarter inflation forecasts as

$$100\left[\left(\frac{PGDP3}{PGDP2}\right)^4 - 1\right].$$

The approach follows the standard practice in the Survey of Professional Forecasters. To calculate individual inflation forecasts, we directly use the equation above. For multiple-quarter-ahead forecasts, we follow the same approach but use multiple-quarter-ahead forecasts of the GDP deflator. For example, to calculate two-quarter-ahead inflation forecasts, we use *PGDP4* and *PGDP3* in the formula above.

For realizations, we use the Real-Time Data Set for Macroeconomists which is also provided by the Philadelphia Fed (link). We use the first-release data for "*Price Index for GNP/GDP (P)*." In 1995Q4, the first-release data for inflation is not available. In this period, we use the second-release data.

To match forecasts and actuals, we align the forecasts to the date for which they were made. For example, the one-quarter ahead forecast made in the 1970Q1 survey is matched with the

		Consensus		Individual	
	Actuals	Forecasts	Errors	Forecasts	Errors
Sample size	210	210	210	7984	7984
Mean	3.36	3.39	-0.03	3.41	0.03
Standard deviation	2.56	2.08	1.44	2.36	1.96
Autocorrelation	0.82	0.97	0.37	0.44	0.16

Table A.1: Summary statistics for one-quarter-ahead GDP deflator inflation forecasts. *Actuals* denotes realized inflation (first-release data). *Errors* equal realized inflation minus forecast. Autocorrelation is measured by the OLS estimate of b in the regression $z_t = a + bz_{t-1} + u_t$; forecaster fixed effects included when calculating the autocorrelation coefficient for individual-level variables.

actual inflation reported for 1970Q2.

In addition, we use data from the FRED database (link) for the Federal Funds Rate (code DFF, link) and the NBER recession indicator (code USRECQ, link). To convert the Federal Funds Rate to a quarterly frequency, we take an average of the daily data.

For the Romer-Romer monetary-surprise measure, we use the extended dataset provided by Wieland and Yang (2020) (link). Wieland and Yang extend the original series of Romer and Romer (2004) up to the end of 2007.

Summary statistics for the final dataset are provided in Table A.1.

C.2 Empirical Mean-Squared Forecast Error Decomposition

In Section 2.2 we demonstrate that the mean-squared forecast error can be decomposed into components due to (i) unconditional bias; (ii) under/overreaction to aggregate shocks; (iii) overreaction to idiosyncratic shocks; and (iv) irreducible error (due to unpredictable shocks). In this section, we discuss the details of our empirical implementation of the decomposition in Section 4.4.1.

Since we are only interested in the combined impact of aggregate shocks, we simplify the decomposition in Eq. (2) by assuming that there is only one aggregate shock present ($J = 1$), as in the model of Section 2.4. Since idiosyncratic shocks wash out at the consensus level, the

mean-squared error of the consensus forecast, $\text{MSE}_{\text{cons.}}$, simplifies to

$$\text{MSE}_{\text{cons.}} = b_{c,0}^2 + \sigma_1^2 \sum_{\ell=1}^{+\infty} [b_{1,\ell}^{(\varepsilon)}]^2 + \sigma_1^2.$$

Here, we denote the unconditional bias of the consensus forecast by $b_{c,0}$. We estimate the unconditional bias, $b_{c,0}$, by the sample average of consensus forecast errors, and we estimate the aggregate bias coefficients by composite bias coefficients. (Recall that composite bias coefficients estimate aggregate bias coefficients consistently when $J = 1$.)

With these estimates at hand, we obtain an estimate of σ_1^2 as

$$\hat{\sigma_1^2} = \frac{\widehat{\text{MSE}}_{\text{cons.}} - \hat{b}_{c,0}^2}{1 + \sum_{\ell=1}^{L} [\hat{b}_{1,\ell}^{(\varepsilon)}]^2},$$

where $\widehat{\text{MSE}}_{\text{cons.}}$ denotes the sample mean-squared forecast error of the consensus forecast, L is the number of composite bias coefficients we estimate, and variables with hats denote sample estimates. We use a truncation parameter of $L = 12$ in our empirical implementation. Note that this approach to estimating σ_1^2 does not require us to take a stand on the true data-generating process for inflation (apart from our simplifying assumption that there is only one aggregate shock). For example, we need not assume that the true data-generating process is an AR(1).

Again applying Eq. (5), the mean-squared error of an individual forecaster i is

$$\text{MSE}_i = b_{i,0}^2 + \sigma_1^2 \sum_{\ell=1}^{+\infty} [b_{1,\ell}^{(\varepsilon)}]^2 + \sigma_\omega^2 \sum_{\ell=1}^{+\infty} [b_\ell^{(\omega)}]^2 + \sigma_1^2. \tag{A.3}$$

In principle, one could estimate this decomposition forecaster-by-forecaster directly. However, many forecasters are in our data for a limited number of periods only. To reduce sampling error, we therefore take the expectation of Eq. (A.3) across forecasters to obtain

$$\mathbb{E}_i[\text{MSE}_i] = \mathbb{E}_i[b_{i,0}^2] + \sigma_1^2 \sum_{\ell=1}^{+\infty} [b_{1,\ell}^{(\varepsilon)}]^2 + \sigma_\omega^2 \sum_{\ell=1}^{+\infty} [b_\ell^{(\omega)}]^2 + \sigma_1^2,$$

where $\mathbb{E}_i[\cdot]$ highlights that the expectation is taken across forecasters. We use the equation above for our empirical decomposition when using forecaster-level data.

As with the consensus forecast, we estimate the unconditional bias $b_{i,0}$ by the sample average of forecast errors made by forecaster i; denote this estimate by $\hat{b}_{i,0}$. We similarly obtain a sample estimate of the mean-squared error of forecaster i, $\widehat{\text{MSE}}_i$. We estimate $\mathbb{E}_i[\text{MSE}_i]$ by $\sum_{i=1}^N \upsilon_i \cdot \widehat{\text{MSE}}_i$ and $\mathbb{E}_i[b_{i,0}^2]$ by $\sum_{i=1}^N \upsilon_i \cdot (\hat{b}_{i,0})^2$, where N is the number of forecasters in the dataset, and υ_i are weights based on the number of forecasts made by forecaster i. Intuitively, we give a higher weight to forecasters that make more forecasts when calculating sample averages. Formally, let $\mathcal{I}_{it}$ denote an indicator variable that equals one if forecaster i made a forecast at time t, and zero otherwise. Furthermore, let $T_i = \sum_{t=1}^T \mathcal{I}_{it}$ denote the total number of periods in which forecaster i made forecasts, where T gives the time-series dimension of the dataset. We calculate sample weights υ_i as $\upsilon_i = T_i / \sum_{j=1}^N T_j$. With these estimates at hand, we back out the contribution of the idiosyncratic component to mean-squared error as

$$\sum_{i=1}^N \upsilon_i \cdot \widehat{\text{MSE}}_i - \sum_{i=1}^N \upsilon_i \cdot (\hat{b}_{i,0})^2 - \hat{\sigma}_1^2 \sum_{\ell=1}^L [\hat{b}_{1,\ell}^{(\varepsilon)}]^2 - \hat{\sigma}_1^2.$$

C.3 Decomposing Coibion-Gorodnichenko Estimates

Our main empirical finding is that inflation expectations underreact to aggregate shocks and overreact to idiosyncratic shocks. We now demonstrate that this feature of expectations is key for understanding the diverging estimates of the Coibion and Gorodnichenko (2015, CG for short) regression at the consensus and individual levels (Bordalo, Gennaioli, Ma and Shleifer, 2020).

Decomposing CG Estimates: Theory. Consider the following three CG-type regressions:

$$\begin{aligned} e_{t+1} &= \alpha + \beta_{\text{CG}}^{\text{macro}} \cdot r_t + u_t \\ e_{i,t+1} &= \alpha + \beta_{\text{CG}}^{\text{micro}} \cdot r_{it} + u_{it} \\ e_{i,t+1} - e_{t+1} &= \alpha + \beta_{\text{CG}}^{\text{idio.}} \cdot (r_{it} - r_t) + u_{it} \end{aligned} \tag{A.4}$$

As before, $e_{i,t+1} = x_{t+1} - \mathbb{F}_{it}[x_{t+1}]$ denotes individual forecast errors, and $r_{it} = \mathbb{F}_{it}[x_{t+1}] - \mathbb{F}_{i,t-1}[x_{t+1}]$ stands for individual forecast revisions; e_t and r_t denote forecast errors and forecast revisions at the consensus level.

The first regression ("macro CG") in Eq. (A.4) is the original regression put forward by Coibion and Gorodnichenko (2015) to estimate information frictions using consensus-level forecasts; see also Nordhaus (1987). The second regression ("micro CG") in Eq. (A.4) is the application of the original CG method to individual-level forecasts, as in Bordalo, Gennaioli, Ma and Shleifer (2020). Finally, the third regression ("idiosyncratic CG") is a novel regression that we propose to decompose the micro CG coefficient into its aggregate and idiosyncratic components.[31] By the Frisch-Waugh-Lovell Theorem, the idiosyncratic CG regression can be estimated by simply adding time fixed effects to the micro CG regression.

Idiosyncratic shocks wash out in the consensus forecast (as the number of forecasters grows large). Hence, the macro CG regression only depends on aggregate bias coefficients. Under the maintained assumption of Section 2 that forecasters respond to aggregate shocks homogeneously,[32] deviations from consensus are only driven by idiosyncratic shocks. Hence, the idiosyncratic CG regression only depends on the idiosyncratic bias coefficients. Finally, the micro CG estimate depends on both aggregate and idiosyncratic bias coefficients.

Suppose that we have access to a balanced panel dataset of observations $\{e_{i,t+1}, r_{it}\}$ for $i = 1, 2, \ldots, I$, and $t = 1, 2, \ldots, T$, where I denotes the number of forecasters, and T is the time-series dimension of the dataset. The next proposition demonstrates that the micro CG

[31] In independent work, Angeletos, Huo and Sastry (2021, Appendix C) also propose a modified version of the micro CG regression and a similar decomposition. Their proposed regression combines the micro and idiosyncratic CG regressions into a single specification. For a particular model of expectation formation, the authors show how this regression can be used to back out the level of information frictions. Our goal, in contrast, is to decompose micro CG estimates into aggregate and idiosyncratic components, without any appeal to a specific model of expectation formation.

[32] The theoretical model of Bordalo, Gennaioli, Ma and Shleifer (2020) implicitly makes this assumption.

	(1)	(2)
$\hat{\beta}_{\text{CG}}^{\text{macro}}$	0.40	–
	(0.40)	
$\hat{\beta}_{\text{CG}}^{\text{micro}}$	-0.35	-0.38
	(0.05)	(0.04)
$\hat{\beta}_{\text{CG}}^{\text{idio.}}$	-0.49	-0.49
	(0.03)	(0.03)
Weight on Idiosyncratic	0.84	0.88
Forecaster FE		✓

Table A.2: Estimates of Coibion and Gorodnichenko (2015, CG) regressions in Eq. (12). $\hat{\beta}_{\text{CG}}^{\text{macro}}$: Estimates of the original CG regression using consensus-level data. $\hat{\beta}_{\text{CG}}^{\text{micro}}$: Estimates of the CG regression using individual-level data. $\hat{\beta}_{\text{CG}}^{\text{idio.}}$: Estimates of the idiosyncratic CG regression (micro CG regression with time fixed effects). *Weight on Idiosyncratic* gives the implied weight that the micro CG coefficient implicitly puts on idiosyncratic shocks, calculated as $(\hat{\beta}_{\text{CG}}^{\text{micro}} - \hat{\beta}_{\text{CG}}^{\text{macro}})/(\hat{\beta}_{\text{CG}}^{\text{idio.}} - \hat{\beta}_{\text{CG}}^{\text{macro}})$. Standard errors in parentheses.

estimate is a weighted average of the macro and idiosyncratic CG coefficients.[33]

Proposition 9. *The OLS estimate of the micro CG coefficient in Eq.* (12) *can be decomposed as*

$$\hat{\beta}_{CG}^{micro} = w \cdot \hat{\beta}_{CG}^{macro} + (1 - w) \cdot \hat{\beta}_{CG}^{idio.},$$

where the weight on the macro CG coefficient is $w \equiv \frac{\sum_{i=1}^{I}\sum_{t=1}^{T}(r_t - \bar{r})^2}{\sum_{i=1}^{I}\sum_{t=1}^{T}(r_{it} - \bar{r})^2}$.

Proof. The OLS estimators of the regression coefficients are given by

$$\hat{\beta}_{\text{CG}}^{\text{micro}} = \frac{\sum_{i=1}^{I}\sum_{t=1}^{T}(e_{i,t+1} - \bar{e})(r_{it} - \bar{r})}{\sum_{i=1}^{I}\sum_{t=1}^{T}(r_{it} - \bar{r})^2}$$
$$\hat{\beta}_{\text{CG}}^{\text{macro}} = \frac{\sum_{t=1}^{T}(e_{t+1} - \bar{e})(r_t - \bar{r})}{\sum_{t=1}^{T}(r_t - \bar{r})^2}$$
$$\hat{\beta}_{\text{CG}}^{\text{idio.}} = \frac{\sum_{i=1}^{I}\sum_{t=1}^{T}(e_{i,t+1} - e_{t+1})(r_{it} - r_t)}{\sum_{i=1}^{I}\sum_{t=1}^{T}(r_{it} - r_t)^2}$$

[33] The result mirrors the classic result in panel-data econometrics that the "pooled OLS" estimator is a matrix weighted average of the "between" and "within" estimators (see, e.g., Baltagi, 2005, Chapter 2.3).

where $\bar{r} = \frac{1}{IT}\sum_{i=1}^{I}\sum_{t=1}^{T} r_{it}$ and $\bar{e} = \frac{1}{IT}\sum_{i=1}^{I}\sum_{t=1}^{T} e_{i,t+1}$. First, calculate that

$$\sum_{i=1}^{I}\sum_{t=1}^{T}(r_{it} - r_t)^2 = \sum_{i=1}^{I}\sum_{t=1}^{T}(r_{it} - \bar{r})^2 - I\sum_{t=1}^{T}(r_t - \bar{r})^2. \tag{A.5}$$

We now notice that

$$\begin{aligned}\sum_{i=1}^{I}\sum_{t=1}^{T}(e_{i,t+1} - e_{t+1})(r_{it} - r_t) = \sum_{i=1}^{I}\sum_{t=1}^{T}(e_{i,t+1} - \bar{e})(r_{it} - \bar{r}) \\ - I\sum_{t=1}^{T}(e_{t+1} - \bar{e})(r_t - \bar{r}).\end{aligned} \tag{A.6}$$

Using Eqs. (A.5) and (A.6) and the definition of w yields the expression stated in the proposition.

□

Empirical Results. In Table A.2, we use Proposition 9 to empirically decompose the micro CG estimates into their aggregate and idiosyncratic components. As in Bordalo, Gennaioli, Ma and Shleifer (2020), we obtain a positive macro CG coefficient ($\hat{\beta}_{\text{CG}}^{\text{macro}} = 0.40$) and a negative micro CG coefficient ($\hat{\beta}_{\text{CG}}^{\text{micro}} = -0.35$ without forecaster fixed effects and $\hat{\beta}_{\text{CG}}^{\text{micro}} = -0.38$ when forecaster fixed effects are included). Point estimates of the idiosyncratic CG regression are $\hat{\beta}_{\text{CG}}^{\text{idio.}} = -0.49$, both with and without forecaster fixed effects. A negative idiosyncratic CG coefficient is consistent with our previous finding of overreaction to idiosyncratic shocks.

Applying Proposition 9, we find that the micro CG regression puts more than 80% of weight on idiosyncratic shocks, and less than 20% to aggregate shocks. Hence, the reason why the micro CG coefficient is negative is because individual expectations overreact to idiosyncratic shocks. The positive macro CG coefficient stems from underreaction to aggregate shocks. As a result, our main empirical finding of underreaction to aggregate shocks and overreaction to idiosyncratic shocks is consistent—and helps explain—the puzzling evidence from micro and macro CG regressions.

C.4 State and Time Dependence

Our methodology can be used to investigate whether biases in expectation formation vary across time and state of the economy, as we now illustrate. Allowing for state and time dependence

is important empirically. For instance, Baker, McElroy and Sheng (2020) document that information rigidity declines substantially in the aftermath of natural disasters.

State Dependence. For a given horizon ℓ, we estimate the following generalization of Eq. (16) using consensus-level forecasts:

$$e_{t+\ell} = \beta_{0,0} + \beta_{0,1} I_t + \beta_{1,0} e_t + \beta_{1,1} I_t e_t + \sum_{k=1}^{K-1} \pi_{k,0} e_{i,t-k} + \sum_{k=1}^{K-1} \pi_{k,1} I_t e_{i,t-k} + u_{t+\ell}. \tag{A.7}$$

Here, $I_t \in \{0, 1\}$ is a dummy variable that equals 1 if the economy is in a particular state at time t (such as a high-inflation state). Following Ramey and Zubairy (2017), we allow all coefficients in the regression to differ across the two states. The estimated composite bias coefficient when the economy is in a particular state ($I_t = 1$) is then given by $\hat{b}_\ell = -(\hat{\beta}_{1,0} + \hat{\beta}_{1,1})$. We test for state dependence by employing a standard t-test for $H_0 : \beta_{1,1} = 0$.

Table A.3 collects the findings for a number of different states of the economy. The estimated bias coefficients indeed exhibit state dependence. Expectations display more underreaction in high-inflation periods, with a difference that is statistically significant. In addition, expectations show less underreaction during recessions, especially at deeper lags, and composite bias coefficients are statistically indistinguishable from zero in recessions. However, state dependence in the estimated bias coefficients is primarily driven by state dependence in the behavior of actual inflation. Table A.4 shows that exactly the same qualitative pattern of state dependence is observed for actual inflation. Hence, the results from this section corroborate the previous evidence of underreaction to aggregate shocks, suggesting that forecasters are slow to react to state dependence in actual inflation.

Time Dependence. We can also perform a similar exercise for time variation in biases. To do so, we estimate composite bias coefficients using a rolling-window regression. The results are given in Figure A.2 which plots the first composite bias coefficient over time. Recall that the first composite bias coefficient provies a summary measure of how expectations react to current news. There is substantial time variation in the estimated bias coefficient. Once again, however, this time variation is driven by time variation in the persistence of actual inflation, as shown in Figure A.3.

	Baseline	High Inflation	Recession	Great Moderation	Large FE
Lag 1	-0.31 (0.10)	-0.54 (0.14) [-2.35]	-0.02 (0.31) [1.25]	-0.30 (0.09) [-0.02]	-0.25 (0.11) [0.84]
Lag 2	-0.08 (0.09)	-0.07 (0.13) [-0.11]	-0.11 (0.14) [-0.08]	0.10 (0.09) [1.67]	-0.14 (0.13) [-0.39]
Lag 3	-0.27 (0.11)	-0.34 (0.21) [-0.61]	0.01 (0.17) [1.88]	-0.15 (0.12) [0.89]	-0.45 (0.13) [-2.36]
Lag 4	-0.22 (0.10)	-0.49 (0.13) [-2.69]	0.26 (0.20) [3.09]	-0.19 (0.10) [0.25]	-0.38 (0.15) [-1.05]

Table A.3: State dependence in composite bias coefficients. Standard errors in parentheses; t-statistics of a test of no state dependence ($\beta_{1,1} = 0$ in Eq. (A.7)) in brackets. High absolute values of the t-statistics are evidence of state dependence, with negative t-statistics indicating more underreaction.

High Inflation: Realized inflation above the 75$^{\text{th}}$ percentile of sample values. *Recession*: The economy is in a recession as given by the NBER recession indicator. *Great Moderation*: Dummy variable that is equal to 1 between 1985-01-01 and 2006-12-31. *Large FE*: Forecast error above the 75$^{\text{th}}$ percentile of sample values.

C.5 Bootstrap Standard Errors

In this section we discuss how we calculate bootstrap standard errors for the parameter estimates in Section 4.4.2.

To account for time-series dependence, we use a moving-blocks bootstrap (Kunsch, 1989; Gonçalves, 2011) with a block size of 20. Since we employ quarterly data, our choice implies a block size of five years. As shown in Section 4, most bias coefficients are close to zero already two years after the initial shock. Hence, it seems unlikely that there is any substantial time-series dependence in forecast errors after five years.

More specifically, for each bootstrap replication $b = 1, 2, \ldots, B$, we:

(1) Create a forecaster-level bootstrap dataset by drawing blocks with replacement in the

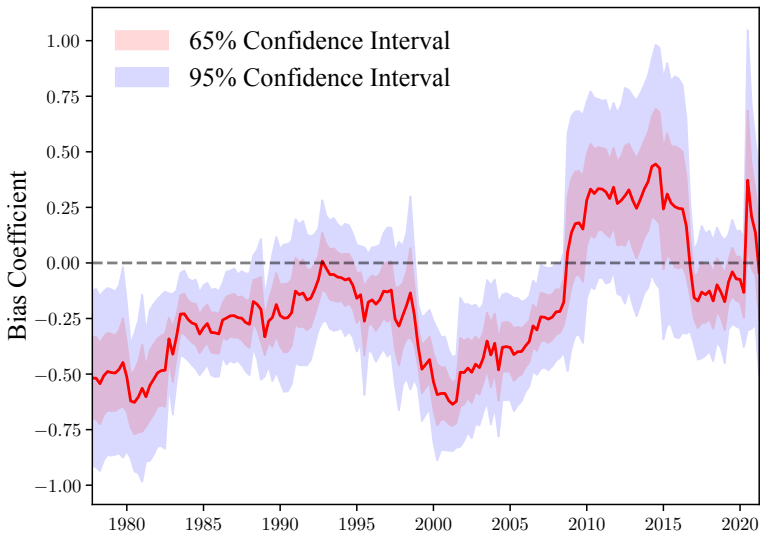

Figure A.2: Rolling window estimates of the first composite bias coefficient, $\hat{b}_1$. A window of 32 quarters (8 years) is used for estimation.

	Baseline	High Inflation	Recession	Great Moderation	Large FE
Lag 1	0.46 (0.10)	0.64 (0.14) [2.46]	0.32 (0.26) [-0.79]	0.36 (0.09) [-0.77]	0.39 (0.11) [-0.80]
Lag 2	0.41 (0.10)	0.33 (0.17) [0.59]	0.46 (0.17) [0.29]	0.07 (0.12) [-2.51]	0.41 (0.13) [0.14]
Lag 3	0.50 (0.12)	0.48 (0.21) [1.01]	0.25 (0.26) [-1.56]	0.30 (0.08) [-1.44]	0.64 (0.16) [1.66]
Lag 4	0.57 (0.14)	0.64 (0.17) [2.11]	0.07 (0.19) [-4.20]	0.43 (0.09) [-0.85]	0.64 (0.20) [0.32]

Table A.4: State dependence in the univariate IRF of quarterly GDP deflator inflation. Standard errors in parentheses; t-statistics of a test of no state dependence in brackets. High absolute values of the t-statistics indicate evidence of state dependence, with negative t-statistics indicating less persistence. Definitions of the different states are provided in Table A.3.

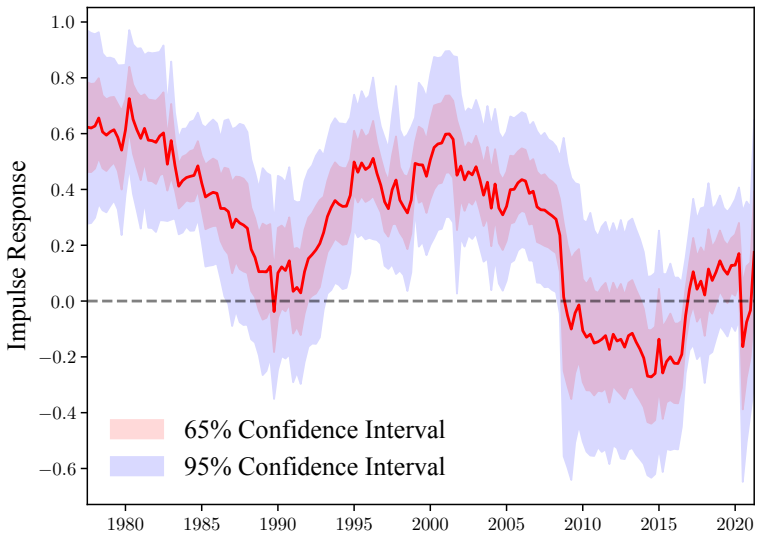

Figure A.3: Rolling window estimates of the univariate IRF of quarterly GDP deflator inflation (first lag). The estimation uses local projections and a window of 32 quarters (8 years); Newey-West standard errors with 4 lags are used to calculate the confidence intervals.

time series dimension. In particular, denoting the time-series dimension of the original dataset by T, we construct a list of $M \equiv \lceil T/20 \rceil$ integers $L^{(b)} = \{t_1^{(b)}, t_2^{(b)}, \ldots, t_M^{(b)}\}$ by drawing uniformly with replacement from $\{1, 2, \ldots, T-19\}$. We then construct the bootstrap dataset by iterating through the list $L^{(b)}$ and stacking observations. That is, for $j = 1, 2, \ldots, M$, we:

- Select all observations from the original dataset whose time index t satisfies $t \geq t_j^b$ and $t \leq t_j^b + 19$;
- Save these observations in the bootstrap dataset.

(2) Create a consensus-level bootstrap dataset by taking medians of individual-level inflation forecasts. Note that because we resample blocks, we work with inflation forecasts (rather than price-level forecasts) in order not to introduce discontinuities in the bootstrap datasets;

(3) Estimate idiosyncratic bias coefficients using the forecaster-level bootstrap dataset (as in Section 4);

(4) Estimate composite bias coefficients using the consensus-level bootstrap dataset, (as in Section 4);

(5) Fit the models of Section 4.4.2 by matching the bootstrap empirical bias coefficients from steps (3) and (4).

We estimate standard errors by the standard deviation of the bootstrap estimates in step (5).

For the noisy-diagnostic model, we modify the above algorithm slightly. As discussed in Section 4.4.2, the noisy-diagnostic model is not well identified by the empirical bias coefficients. In particular, there are many values of the diagnosticity parameter θ and Kalman gain G that yield a similar level of fit. As a result, in some of the bootstrap replications, we obtain extreme values for the estimated diagnosticity parameter. To avoid this problem affecting the standard-error estimates, we impose an ad-hoc constraint that $\theta \leq 10$, and calculate the standard error only from bootstrap realizations that satisfy this constraint.

In unreported results, we have also experimented with an alternative method for calculating standard errors. We first estimated AR(4) models on forecast errors (in the case of composite bias coefficients) or deviations from consensus (in the case of idiosyncratic bias coefficients) using the original dataset. We then used the point estimates and the estimated variance-covariance matrix of the AR(4) coefficients to simulate bias coefficients. This approach yielded similar standard errors to the moving-blocks bootstrap. This alternative approach, however, implicitly assumes that estimated composite and idiosyncratic coefficients are independent, an assumption that is likely violated in the data. On the other hand, this alternative approach is substantially cheaper computationally and may therefore be appealing in certain empirical settings.

To statistically test whether the difference in goodness-of-fit between two models is statistically significant, we used the following approach. For each bootstrap simulation, we calculate the difference in the sum of squared residuals (SSR) between the two models in question and denote it by ΔSSR. We then estimate the standard error of ΔSSR by the standard deviation of the bootstrap realizations. The difference in the goodness-of-fit on the original sample can then be tested using a simple t-test. Figure A.4 plots the histogram of bootstrap realizations of ΔSSR for the two sets of models that have the same number of free parameters. We observe that ΔSSR exhibits substantial variation in both cases, suggesting that the difference in the goodness-of-fit

observed in the original dataset is not statistically significant.

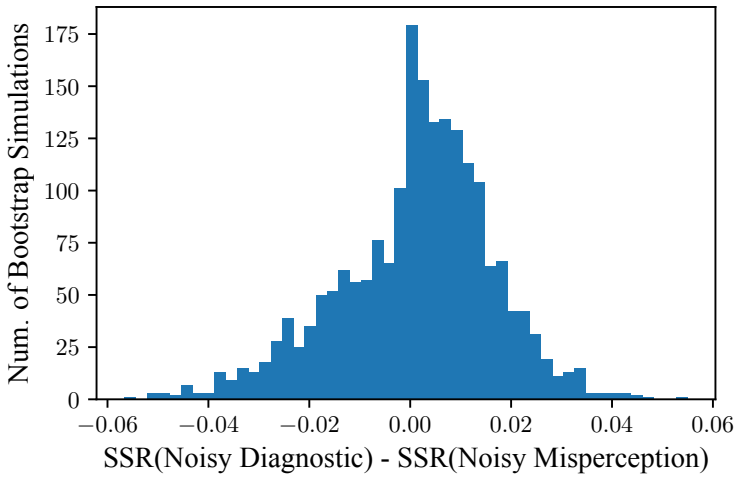

(a) ΔSSR: Noisy-Diagnostic Model vs. Noisy-Misperception Model

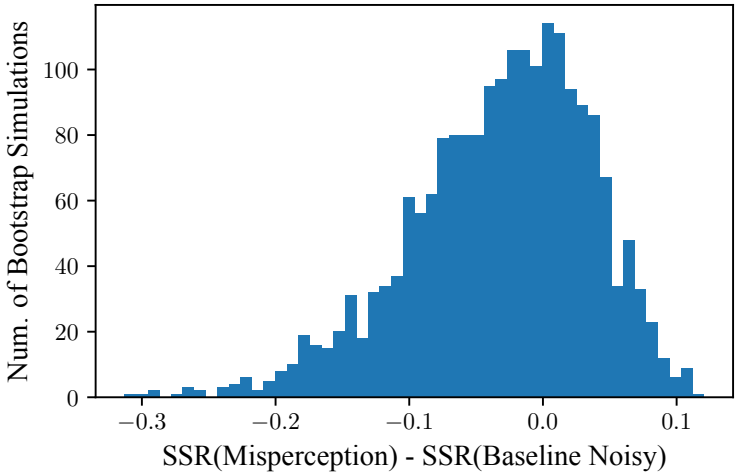

(b) ΔSSR: Misperception Model vs. Baseline-Noisy Model

Figure A.4: Histograms of bootstrap realizations of ΔSSR, the difference in goodness-of-fit between two models. Moving-blocks bootstrap with 2,000 bootstrap replications. *Top panel*: Difference in goodness-of-fit between the noisy-diagnostic and noisy-misperception models. *Bottom panel*: Difference in goodness-of-fit between the misperception and baseline-noisy models.

C.6 Additional Empirical Results

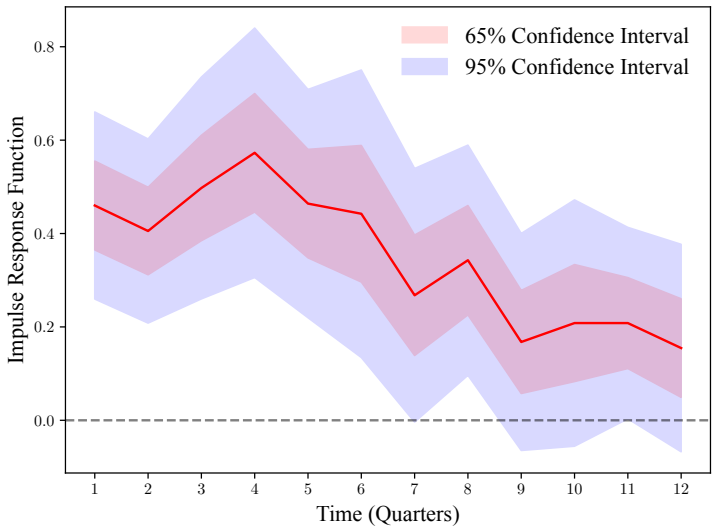

Figure A.5: Univariate impulse response function of quarterly GDP deflator inflation. The estimation uses local projections; Newey-West standard errors with $\max\{5, \ell - 1\}$ lags are used to calculate the confidence intervals where $\ell = 1, 2, \ldots, 12$ denotes the horizon of the local projection.

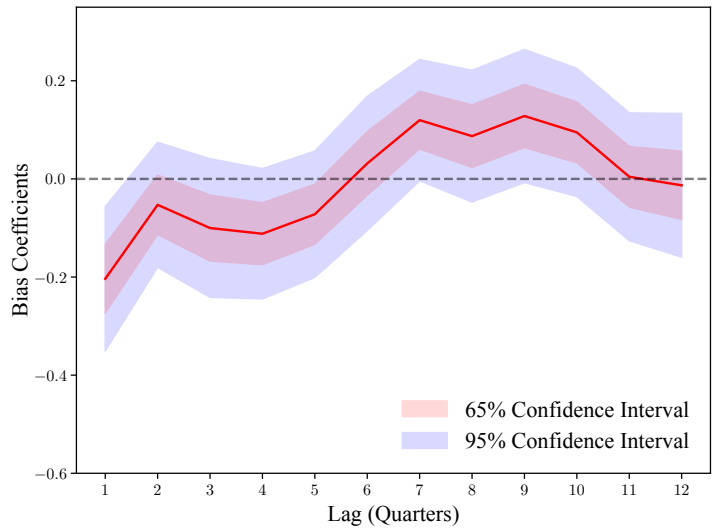

Figure A.6: Composite bias coefficients: Individual forecasts with winsorization at 5% and 95%. Same regression specification as in Figure **??** but with forecast errors winsorized at 5% and 95% (using full sample percentiles).

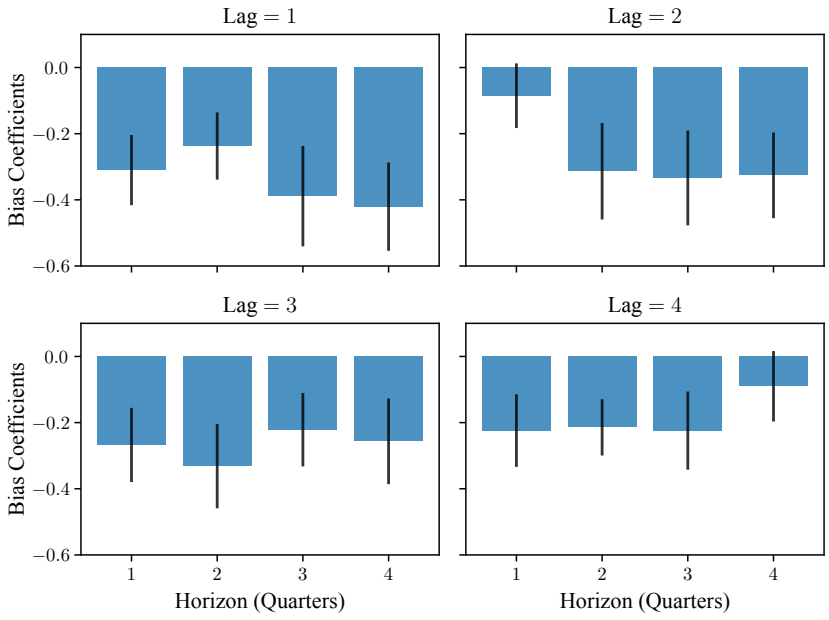

Figure A.7: Composite bias coefficients for one-, two-, three-, and four-quarter ahead forecasts. Forecast horizons are given by the *x*-axis of each subplot. "Lag = 1" gives the first bias coefficient (i.e., response to current news), while "Lag = 4" gives the fourth bias coefficient (i.e., response to news that arrived three quarters ago). Plus/minus one standard error is shown by the error bars. The regression specification is the same as in Figure **??**.

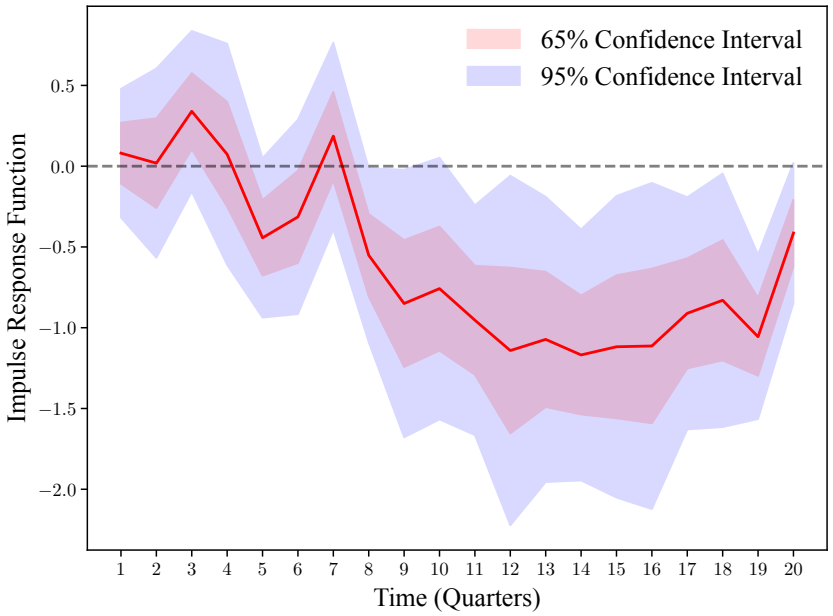

Figure A.8: Response of inflation to monetary policy shocks. The response is estimated by $\hat{\delta}_1^{(\ell)}$ in the regression

$$x_{t+\ell} = \alpha^{(\ell)} + \sum_{s=1}^{4} \delta_s^{(\ell)} \text{FFR}_{t+1-s} + \sum_{s=1}^{4} \beta_s^{(\ell)} x_{t+1-s} + u_{t+\ell},$$

for $\ell = 1, 2, \ldots, 20$, instrumenting the Federal Funds Rate, FFR, at time t with the Romer and Romer (2004) monetary policy shock. Newey-West standard errors with $\max\{5, \ell - 1\}$ lags are used to calculate the confidence intervals where ℓ denotes the horizon of the local projection.

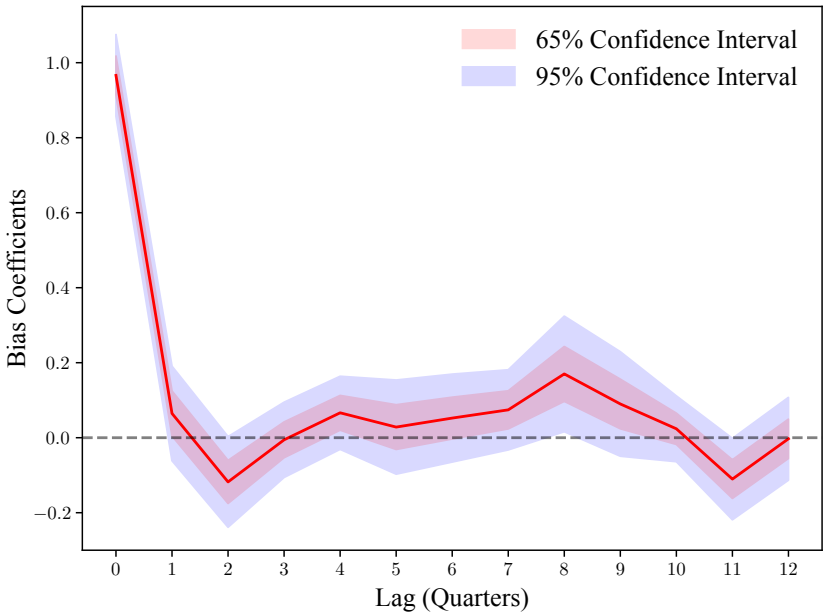

Figure A.9: Bias coefficients w.r.t. to idiosyncratic shocks. Individual-level forecasts are used in the IV procedure of Section 3.4, instrumenting $\mathbb{F}_{it}[x_{t+1}]$ with idiosyncratic shocks estimated using the single-factor procedure (with $K = 3$) in Juodis and Kucinskas (2019); standard errors clustered by both forecaster and quarter.

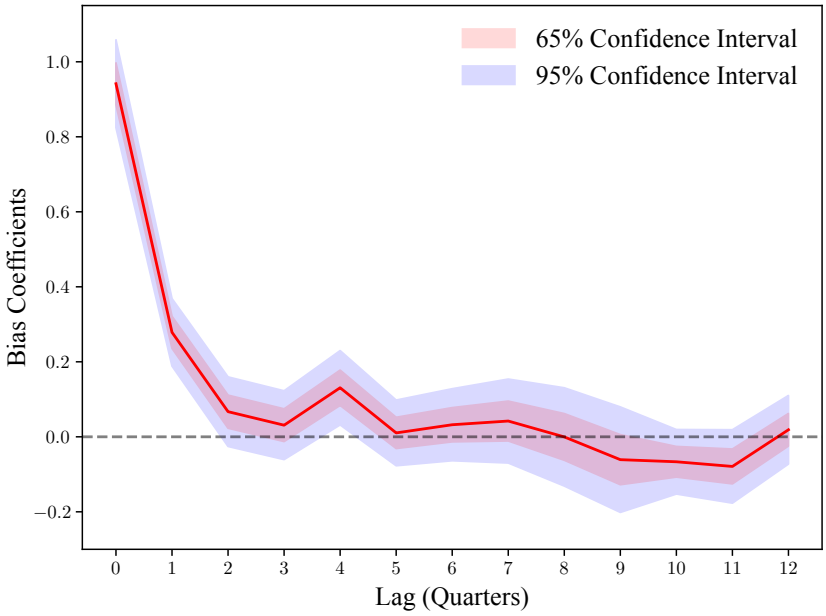

Figure A.10: Bias coefficients w.r.t. to idiosyncratic shocks with winsorization at 5% and 95%. Same regression specification as in Figure **??** but with variables winsorized at 5% and 95% (using full sample percentiles).

Appendix D Additional Proofs

D.1 Bordalo, Gennailoli, Ma, and Shleifer (2020), Proposition 1, with misperception

In Bordalo, Gennaioli, Ma and Shleifer (2020) (BGMS), Eq. (11) is obtained from a micro-founded model of expectation formation in which all parameter values are perceived correctly (Proposition 1 in BGMS). Fortunately, Proposition 1 in BGMS remains valid with misperception, provided that representativeness is defined using perceived density functions. Using the notation in BGMS, define representativeness as $R_t(x) \equiv \hat{f}(x|S_t^i)/\hat{f}(x|S_{t-1}^i \cup \{x_{t|t-1}^i\})$, where $\hat{f}(x|\cdot)$ is a density evaluated using the perceived parameter values. Then, define the diagnostic posterior density function as $f^\theta(x_t|S_t^i) \equiv \hat{f}(x_t|S_t^i)R_t(x_t)^\theta Z_t^{-1}$, where Z_t is a normalization factor. It is straightforward to check that the proof of Proposition 1 in BGMS goes through without any modifications, and it yields Eq. (11).

References: Online Appendix

Angeletos, George-Marios, Zhen Huo, and Karthik A. Sastry, "Imperfect Macroeconomic Expectations: Evidence and Theory," *NBER Macroeconomics Annual*, May 2021, *35* (1), 1–86.

Augenblick, Ned and Eben Lazarus, "Restrictions on Asset-Price Movements Under Rational Expectations: Theory and Evidence," Working paper 2018.

_ **and Matthew Rabin**, "Belief Movement, Uncertainty Reduction, and Rational Updating," *The Quarterly Journal of Economics*, 2021, *136* (2), 933–985.

Baker, Scott R, Tucker S. McElroy, and Xuguang S Sheng, "Expectation Formation Following Large, Unexpected Shocks," *The Review of Economics and Statistics*, may 2020, *102* (2), 287–303.

Baltagi, Badi H., *Econometric Analysis of Panel Data*, John Wiley and Sons, 2005.

Ben-David, Itzhak, John R. Graham, and Campbell R. Harvey, "Managerial Miscalibration," *The Quarterly Journal of Economics*, 2013, *128* (4), 1547–1584.

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

Cochrane, John, *Asset Pricing: Revised Edition*, Princeton: Princeton University Press, 2005.

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

da Silveira, Rava Azeredo and Michael Woodford, "Noisy Memory and Over-Reaction To News," *American Economic Review: Papers & Proceedings*, 2019, *109*, 557–561.

de Silva, Tim and David Thesmar, "Noise in Expectations: Evidence from Analyst Forecasts," Working paper 2021.

Diebold, Francis X. and Jose A. Lopez, "Forecast Evaluation and Combination," in G.S. Maddala and C.R. Rao, eds., *Handbook of Statistics (Statistical Methods in Finance)*, Vol. 14, Amsterdam: Elsevier, 1996, chapter 8, pp. 241–268.

Dominitz, Jeff and Charles F. Manski, "Measuring and Interpreting Expectations of Equity Returns," *Journal of Applied Econometrics*, 2011, *26*, 352–370.

Drerup, Tilman, Benjamin Enke, and Hans Martin von Gaudecker, "The Precision of Subjective Data and the Explanatory Power of Economic Models," *Journal of Econometrics*, 2017, *200* (2), 378–389.

Federal Reserve Bank of Philadelphia, "Survey of Professional Forecasters: Documentation," Link to website 2017.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus, "Five Facts About Beliefs and Portfolios," *American Economic Review*, may 2021, *111* (5), 1481–1522.

Gonçalves, Sílvia, "The Moving Blocks Bootstrap for Panel Linear Regression Models with Indi-

vidual Fixed Effects," *Econometric Theory*, 2011, *27* (5), 1048–1082.

Graham, John R and Campbell R Harvey, "Expectations of Equity Risk Premia, Volatility and Asymmetry From a Corporate Finance Perspective," Working paper 2001.

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

He, Simin and Simas Kucinskas, "Expectation Formation with Correlated Variables," Working paper 2020.

Juodis, Artūras and Simas Kucinskas, "Quantifying Noise," Working paper 2019.

Koop, Gary, M.Hashem Pesaran, and Simon M. Potter, "Impulse Response Analysis in Nonlinear Multivariate Models," *Journal of Econometrics*, 1996, *74* (1), 119–147.

Kunsch, Hans R., "The Jackknife and the Bootstrap for General Stationary Observations," *The Annals of Statistics*, 1989, *17* (3), 1217–1241.

Manski, Charles F., "Measuring Expectations," *Econometrica*, 2004, *72* (5), 1329–1376.

Nordhaus, William D., "Forecasting Efficiency: Concepts and Applications," *The Review of Economics and Statistics*, 1987, *69* (4), 667–674.

Patton, Andrew J. and Allan Timmermann, "Properties of Optimal Forecasts Under Asymmetric Loss and Nonlinearity," *Journal of Econometrics*, 2007, *140* (2), 884–918.

Ramey, Valerie and Sarah Zubairy, "Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data," *Journal of Political Economy*, 2017, *126* (2), 696277.

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

Vissing-Jorgensen, Annette, "Perspectives on Behavioral Finance: Does 'Irrationality' Disappear with Wealth? Evidence from Expectations and Actions," *NBER Macroeconomics Annual*, 2004, *18*, 139–208.

Wieland, Johannes F. and Mu-Jeung Yang, "Financial Dampening," *Journal of Money, Credit and Banking*, 2020, *52* (1), 79–113.