Smooth Diagnostic Expectations*

Francesco Bianchi
Johns Hopkins University
CEPR and NBER

Cosmin Ilut
Duke University
NBER

Hikaru Saijo
UC Santa Cruz

Abstract

We show that the formalization of representativeness (Kahneman and Tversky (1972)) developed by Gennaioli and Shleifer (2010) features an intrinsic connection between *uncertainty* and *overreaction*. In the time series domain, we develop this connection in a *smooth* version of Diagnostic Expectations (DE), where overreaction varies smoothly with uncertainty. Intuitively, under smoothness, lower uncertainty leaves less room for representativeness to distort beliefs. Smooth DE provides a joint, parsimonious micro-foundation for key features of survey data: overreaction to news, stronger overreaction at longer horizons, overconfidence in subjective uncertainty, and a new stylized fact documented here—overreaction intensifies with higher uncertainty. An analytical Real Business Cycle model with Smooth DE accounts for survey overreaction and overconfidence, as well as three salient business cycle properties: asymmetry, countercyclical macro volatility, and countercyclical micro volatility.

JEL Codes: D91, E32, E71.

Keywords: Diagnostic expectations, uncertainty, learning, overreaction, overconfidence

*Emails: francesco.bianchi@jhu.edu, cosmin.ilut@duke.edu, and hsaijo@ucsc.edu. We are grateful to the Editor and the three anonymous referees for their helpful comments. We thank Nick Bloom, Yuriy Gorodnichenko, Joel Flynn, Chen Lian, Jean-Paul L'Huillier, Karthik Sastry, Andrei Shleifer, Sanjay Singh, Simon Sheng, Alireza Tahbaz-Salehi, Stephen Terry, Alonso Villacorta, as well as conference participants at the ASSA 2024 Annual Meeting, the 2024 NBER Summer Institute, the 2024 Stanford Institute of Theoretical Economics conference, the 2025 Econometric Society World Congress, the Dallas Fed, the Minneapolis Fed, the Philadelphia Fed, Indiana, Notre Dame, Purdue, Wake Forest, Bank of Canada, Bank of Italy, Norges Bank, Lancaster, Bundesbank, Cleveland Fed, and the Canon Institute for Global Studies Conference on Macroeconomic Theory and Policy. The authors gratefully acknowledge the financial support of the NSF (award number 2417397). Saijo gratefully acknowledges the financial support of the Grants-in-Aid for Scientific Research (JP21H04397) from the Japan Society for the Promotion of Science.

1 Introduction

There has been a growing interest in psychological foundations that enrich models of belief formation in Economics. A prominent example is the paradigm of Diagnostic Expectations (DE), developed by Gennaioli and Shleifer (2010) and Bordalo et al. (2016), that formalizes the widely-documented "representativeness heuristic" of Kahneman and Tversky (1972). According to this heuristic, agents' subjective forecasts overreact to *new information*, as measured with respect to a reference distribution based on past data, because memory selectively recalls more vividly past events that are more associated with, or representative of, that current news.[1]

In this paper, we emphasize and characterize the deep, pervasive relationship between representativeness and uncertainty, and examine its theoretical and empirical implications. Specifically, we show that the distortion in beliefs, measured as the distance between the subjective distribution and its objective counterpart, is shaped by two properties of the incoming information. First, through a *news channel*, if the new information induces a stronger *revision* in the current objective distribution relative to the reference distribution, the resulting overreaction leads to a larger distortion. Second, through an *uncertainty channel*, if the *same* new information reduces the uncertainty of the current objective distribution, the representativeness distortion weakens. Intuitively, when conditional uncertainty is low, the agent is well informed, leaving less room for representativeness to distort forecasts via selective memory. Conversely, if the same information increases conditional uncertainty, memory plays a larger role, amplifying the distortion in beliefs.

We first illustrate the connection between representativeness and uncertainty in a static, two-state categorical example. We then build on the insights of the static environment to argue that in the time series domain this connection manifests itself in a "smoothed" version of DE (Smooth DE), where overreaction varies smoothly with uncertainty. Finally, we study how Smooth DE provides a joint and parsimonious micro-foundation for key properties of survey data and of the business cycle.

The key to connecting representativeness to the time-series domain is that, under the Smooth DE framework, agents overreact to new information, defined as the difference between the current and a reference information set. Since new information typically alters not only the conditional mean but also *conditional uncertainty*, changes in uncertainty surrounding current and past beliefs shape the extent of the DE distortion. Unlike standard DE models, Smooth DE implies that both overreaction and confidence depend on condi-

[1] Bordalo et al. (2022) provides an overview of DE modeling and related evidence. See Taubinsky et al. (2024) for recent evidence on overreaction in agents' subjective forecasts, uncovering a key role for selective and associative memory recall.

tional uncertainty. As a result, Smooth DE bridges two influential branches of economics that have largely evolved in parallel: the Diagnostic Expectations literature and the Uncertainty literature (Bloom (2009, 2014), Fernández-Villaverde et al. (2015), Basu and Bundick (2017), Bloom et al. (2018), Baker et al. (2024)). It also contributes to a growing literature embedding DE in micro-founded macroeconomic models (Bianchi et al. (2024), L'Huillier et al. (2024), Bordalo et al. (2026), Na and Yoo (2025)).

Representativeness and Uncertainty. Under representativeness, events that become more likely in light of new information receive disproportional weight in subjective probabilistic assessments. We use a classic two-state categorical inference example to highlight that the strength of overreaction under the representativeness heuristic is not monotonic in the revision induced by new information, but rather hump-shaped. Suppose there are only two categories—*red* and *non-red* hair color—and that, prior to receiving any information, red hair has a low probability. As new information (e.g., Irish nationality) increases the likelihood of red hair, overreaction initially rises with the conditional probability assigned to red. This occurs because the *news effect* and the *uncertainty effect* reinforce each other: red becomes more likely, and uncertainty increases as a previously unlikely outcome becomes plausible but not certain. However, once the objective probability crosses the point of maximum uncertainty (50:50 in a two-state distribution), the uncertainty effect begins to offset the news effect. As uncertainty declines, the distortion in beliefs diminishes. In the limit, when the probability of red hair approaches 1, uncertainty vanishes, the uncertainty channel dominates, and subjective and objective distributions converge.

Smooth Diagnostic Expectations. The natural extension of representativeness to the time-series domain involves defining new information as the difference between the information set conditional on current data and a reference set conditional on past data. Under Smooth DE, we show that conditional uncertainty shapes belief distortions in a way analogous to the categorical case. Specifically, when both the current and reference distributions are Normal, we analytically uncover two key properties of Smooth DE that are novel to the Diagnostic Expectations literature.

First, Smooth DE delivers a disconnect between the *objective* and *subjective* level of uncertainty. This is because under Smooth DE, not only the mean, but also the variance of the DE distribution is distorted. Even when new information changes only conditional uncertainty, Smooth DE agents overreact to that information, consistent with the basic representativeness heuristic. For example, when information reduces uncertainty with respect to the reference distribution, agents inflate the probability of events close to the mean, while downplaying the probability of tail events that have become less likely. The result is a distorted subjective distribution that is narrower than the current objective one, implying

overconfidence. Given that typically events close in the future are easier to predict than events far into the future, agents' distorted beliefs will typically feature such overconfidence. However, the Smooth DE framework also predicts *underconfidence* following an increase in conditional uncertainty, like in response to an uncertainty shock (Bloom (2009)).

Second, the same underlying principle of overreaction to new information implies that, under Smooth DE, overreaction to news that shifts the conditional mean increases smoothly with uncertainty. For example, suppose the new information shifts the current density to the right of the reference density. For a given mean shift, the conditional variance influences the relative probability of tail events. When uncertainty is low, overreaction is weak because events near the new mean become more *representative*. In contrast, when uncertainty is high, tail events become more likely, amplifying overreaction. In the limit, as uncertainty approaches zero, overreaction vanishes since uncertainty is fully resolved.

Comparison with standard DE. Smooth DE arises from a minimal yet conceptually important modification to the baseline DE framework developed by Bordalo et al. (2018) (BGS), and aligns closely with the "representativeness heuristic" of Kahneman and Tversky (1972). In the BGS formulation, the reference distribution is centered on the conditional mean under the true density formed at a past time, but shares the same level of uncertainty as the current true distribution. In contrast, since we define new information as the change in information sets, our reference distribution conditions solely on past data and thus reflects the uncertainty prevailing at that time. This distinction in how uncertainty is treated across information sets underpins our novel interaction between representativeness and uncertainty.

Smooth DE is built on the key informational difference between conditional (posterior) and unconditional (prior) distribution. In this sense, our approach relates to Bordalo et al. (2021), which features a sampling by similarity framework that under conditional probability assessments yields a result reminiscent of DE. In that setting, the entire prior distribution plays a major role in memory interference. The shape of the prior distribution also affects the degree of overreaction in the finite-state framework and evidence presented in Ba et al. (2022). In that work, overreaction to signals decreases with their informativeness, consistent with the logic of Smooth DE. Closest to our setting, a version of what we label Smooth DE appears in chapter 5 of Gennaioli and Shleifer (2018). However, despite this early appearance, the growing DE literature has focused on the simplified version proposed in BGS.

Survey evidence. We use the Survey of Professional Forecasters (SPF) to provide new evidence and test the Smooth DE's basic implication that overreaction of forecasts to news is stronger in times of high uncertainty. In doing so, we follow Bordalo et al. (2020) and study the correlation between forecast revisions and the subsequent forecast errors. Our key novel empirical finding supports the basic Smooth DE prediction: at the individual level the

overreaction to news is indeed stronger in periods when uncertainty is higher. To connect with these novel findings, we then extend the signal-extraction setting in Bordalo et al. (2020) to allow for time-varying uncertainty within our Smooth DE framework.

Smooth DE also provides a unified explanation for other key survey facts. First, overreaction to news increases with the forecast horizon (Bordalo et al. (2019), d'Arienzo (2020), Bordalo et al. (2020), Augenblick et al. (2025), Bordalo et al. (2023) and Halperin and Mazlish (2025)). Under Smooth DE, the same piece of information reduces uncertainty by less for longer horizons, leading to stronger overreaction. Second, recent work documents that in surveys firms (i) *overreact* to news and (ii) are *overconfident* in their subjective forecasts (Barrero (2022), Born et al. (2025), Altig et al. (2022) and Born et al. (2022)). While the baseline DE model can account for overreaction, it is silent on overconfidence. Smooth DE can account for both properties since it distorts both the mean and the variance of agents' expectations in a way to typically generate both overreaction and overconfidence.

Business cycle implications. We leverage our theoretical insights to study a parsimonious business cycle model with *time-varying uncertainty*, illustrating how state-dependent overreaction under Smooth DE generates important cyclical implications. We consider an island economy subject to both aggregate and island-specific TFP shocks. Following Bloom et al. (2018), we assume that the volatility of island-specific TFP shocks varies over time and is negatively correlated with aggregate TFP innovations, while the volatility of aggregate TFP remains constant. This simple framework accounts for the survey evidence on overreaction and overconfidence documented by Barrero (2022), and reproduces three key features of the business cycle: (1) *asymmetry*—recessions are sharper than expansions; (2) *countercyclical macro volatility*—time-series variances of macroeconomic variables increase in recessions, despite homoskedastic aggregate TFP shocks; and (3) *countercyclical micro volatility*—cross-sectional variances of microeconomic variables rise during recessions.[2]

First, consider the asymmetry property. A negative economy-wide TFP shock generates higher uncertainty. Hence, agents overreact to the economy-wide TFP shock more than usual, leading to a sharper fall in hours, consumption, and output. In contrast, a positive TFP shock reduces agents' uncertainty and thus overreaction, and the rise in economic activity is milder. The enhanced overreaction to aggregate shocks during recessions also generates the second property, countercyclical macro volatility. The aggregate volatility rises in recessions even when there is no change in the volatility of economy-wide shocks. Third, consider

[2]These properties are well documented in the literature. For example, Neftci (1984), Hamilton (1989), McKay and Reis (2008), and Morley and Piger (2012) show macroeconomic asymmetries using various econometric approaches. Bloom (2009), Fernández-Villaverde et al. (2011), Ilut et al. (2018), Jurado et al. (2015), Basu and Bundick (2017), and Bloom et al. (2018) document that volatility or uncertainty rises in recessions at both micro and macro levels.

countercyclical micro volatility. In recessions, agents face higher uncertainty, so they overreact to the island-specific TFP and, as a result, the cross-sectional variances of island-level hours, output, and consumption increase. Conversely, during expansions, agents' overreactions are milder, and hence the cross-sectional dispersion decreases. Finally, a novel policy implication emerges: a redistributive policy that insures agents against idiosyncratic risk can enhance macroeconomic stability by directly dampening the behavioral amplification.

These mechanisms highlight that micro-level uncertainty and macroeconomic volatility are tightly linked through state-dependent overreaction, which acts as a powerful propagation channel. Our mechanism is thus in line with the theoretical and empirical focus on asymmetric propagation from otherwise symmetric and homoskedastic aggregate shocks, as for example in Ilut et al. (2018) and Straub and Ulbricht (2024), while complementing alternative channels that link high volatility and contractions through left-skewed fundamental shocks, such as in Salgado et al. (2019) and Berger et al. (2020).

2 Representativeness and uncertainty

The building block for Smooth DE is the concept of representativeness. We first consider a simple probability environment that is static and discrete, which also helps to highlight how *uncertainty* shapes the effective manifestation of representativeness.

Representativeness. Consider a discrete probability space (Φ, P), with two random variables: x, the trait that the agent seeks to assess, and D, the available data. The agent looks to form the conditional probability of a given trait $\widehat{x} \in x$, given a particular realized data $d \in D$. A Bayesian agent would simply use the conditional probability $P(\widehat{x}|d) = P(\widehat{x} \cap d)/P(d)$. The environment here is static, without a sense of repeated accumulation of information.

In assessing conditional probabilities, the agent subject to the representativeness heuristic distorts the Bayesian belief. We follow the formalization of representativeness in Gennaioli and Shleifer (2010), Bordalo et al. (2016, 2021). This work explains in detail how representativeness captures the tendency to overweight representative traits, arising due to limited memory and the fact that representative traits are easier to recall. Formally, the representativeness of a given trait $\widehat{x}$ conditional on a particular data realization d compared to reference data, d^{ref}, is defined as the relative frequency of that trait across the two data groups:

$$rep(\widehat{x}|d, d^{ref}) \equiv \frac{P(\widehat{x}|d)}{P(\widehat{x}|d^{ref})} \tag{1}$$

The conditional belief under representativeness is affected by a weight that distorts the

underlying absolute frequency of $\widehat{x}$ conditional on the data d

$$P^{\theta}(\widehat{x}|d, d^{ref}) = P(\widehat{x}|d)weight(\widehat{x}|d, d^{ref}). \tag{2}$$

The weight reflects the effect of the relative frequency in equation (1), and is given by

$$weight(\widehat{x}|d, d^{ref}) = \left[rep(\widehat{x}|d, d^{ref})\right]^{\theta} / Z(d, d^{ref}). \tag{3}$$

with $Z(d, d^{ref})$ a constant of integration such that the distribution integrates to one:

$$Z(d, d^{ref}) = \sum_{\widehat{x} \in x} P(\widehat{x}|d) \left[rep(\widehat{x}|d, d^{ref})\right]^{\theta} \tag{4}$$

The parameter $\theta \geq 0$ measures the extent to which the relative frequency across two conditioning data groups affects judgements. When $\theta = 0$, agent's memory retrieval is perfect and beliefs collapse to the standard frictionless model. When $\theta > 0$, memory is limited and agent's judgments are shaped by representativeness.

Illustration. We consider an example based on a two-state categorical distribution. This example allows us to present the results using simple metrics. In Appendix A, we extend the analysis using information-theoretic measures and allowing for more than two states.

Let Φ be the universe of people, and let x be the hair color, which can assume only two values: "red" or "other": $\widehat{x} = R$ and $\widehat{x} = O$. The data $d \in D$ specifies a particular nationality, with one of them serving as the reference $d^{ref} = W$, corresponding to the World (W) population.[3] The agent observes the nationality of an unknown individual and assesses the probability of red hair using $P^{\theta}(\widehat{x} = R|N, W)$, as defined in equation (2).

A key message of this example is that representativeness induces overreaction, and its strength depends on the overall shape of the distributions of the two data groups being compared. Although this property has typically been outside the main focus of recent applied work on representativeness and diagnostic expectations, it will play an important role in our development of Smooth DE. In particular, we emphasize *three qualitative properties*.

1. **Overreaction**. The first and most immediate characteristic of representativeness is that the distorted probability implies overreaction to the new information conveyed by the nationality, consistently with equation (2). We illustrate this property graphically in Figure 1, panel (a.1), where we plot the distorted $P^{\theta}(\widehat{x} = R|N, W)$ for different values of

[3]The particular numbers we use are chosen solely to illustrate key features of representativeness and do not necessarily reflect the actual incidence of hair colors in the world population. In fact, we implicitly assume that there exist nationalities for which the probability of red hair is 1.

Figure 1: News and uncertainty in a two-state discrete distribution

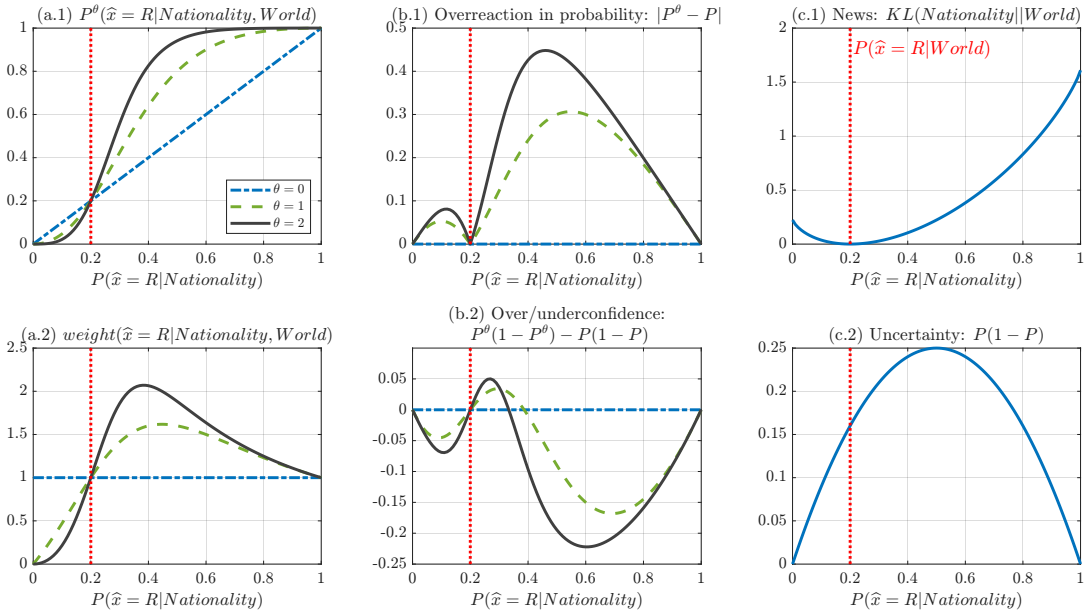

Notes: The figure displays how, as we vary the actual $P(\widehat{x}=R|Nationality)$, the following objects change: (a.1) the distorted probability of red hair, (a.2) the distorting weight on the red hair probability, (b.1) the overreaction in probability, measured by the absolute value of the difference between distorted and objective probabilities, (b.2) overconfidence measured as the difference in uncertainty under the distorted and objective distributions, (c.1) the news component, measured using the KL defined in equation (5), and (c.2) the objective uncertainty. We fix the reference probability of red hair at $P(\widehat{x}=R|W)$=0.2.

$P(\widehat{x}=R|N)$ and diagnosticity parameter values $\theta=1,2$, as well as $\theta=0$ corresponding to no memory distortion. The probability of red hair under the reference data is $P(\widehat{x}=R|W)=0.2$, marked by a vertical red dotted line. When the probability of red hair for a given Nationality is higher (lower) than under the reference data, the distorted probability is higher (lower) than the actual probability for that Nationality. The magnitude of the difference is increasing in θ. This pattern is driven by the behavior of the distorting *weight* $(\widehat{x}=R|N,W)$ defined by equations (3) and (4). We plot this weight in Panel (a.2), showing that it is indeed non-monotonic in the actual frequency: the weight is increasing from an underlying $P(\widehat{x}=R|N)=0$ to an intermediate value between the reference frequency of 0.2 and the upper bound of 1 and then starts decreasing.

Importantly, Panel (a.1) also reveals that the degree of overreaction exhibits a *double hump-shaped* pattern. To better illustrate this key feature, Panel (b.1) plots the *absolute value* of the difference: $\left|P^\theta(\widehat{x}=R\mid N,W)-P(\widehat{x}=R\mid N)\right|$. The overprediction of the *other* hair color peaks at an intermediate value between 0 and the reference point of 0.2, when red hair became objectively less likely. Conversely, the overprediction of *red* hair peaks at an

intermediate value between the reference point of 0.2 and the upper bound of 1.

2. **News and uncertainty.** We now explore the fundamental statistical forces that drive the *double hump-shaped* overreaction of the distorted probability. In particular, we separate two channels that affect the relation between overreaction and the information conveyed by current and reference distributions.

On the one hand, the *weight* $(\widehat{x} = R|N, W)$ in equation (3) features a *news effect*. When the knowledge of the Nationality implies a large *revision* in the frequency of red hair, the representativeness $rep(\widehat{x} = R|N, W)$ in equation (1) changes by more. For a given constant of integration $Z(N, W)$, the distorting weight would thus imply larger overreaction when the revision is larger. On the other hand, there is an *uncertainty effect*: the more the knowledge of a given nationality decreases the uncertainty over the hair color, the less room there is for overreaction to effectively manifest itself. This intuitive force is formally captured by equation (3) through the constant of integration $Z(N, W)$ of equation (4).

The news and uncertainty effects are jointly determined by the shapes of the distributions that get compared through representativeness. The *news* component is large when the distribution implied by knowing the nationality is very different from the reference distribution. We thus measure the size of the *news* component with the Kullback-Leibler (KL) divergence between the current and reference distributions:

$$KL(\mathrm{N}||\mathrm{W}) = P(\widehat{x} = R|N) \ln\left(\frac{P(\widehat{x} = R|N)}{P(\widehat{x} = R|W)}\right) + P(\widehat{x} = O|N) \ln\left(\frac{P(\widehat{x} = O|N)}{P(\widehat{x} = O|W)}\right) \qquad (5)$$

Panel (c.1) shows that the KL divergence is U-shaped. It reaches its minimum value of zero when knowing the nationality does not lead to any revision: $P(\widehat{x} = R \mid N) = P(\widehat{x} = R \mid W)$. It increases progressively, both to the left and to the right of the reference frequency of 0.2, as the new information leads to a departure from that reference. The divergence peaks when the nationality implies that the probability of red hair is 1, as this information is maximally different from the reference probability of 0.2.

To measure *uncertainty* in the objective distribution of hair colors for a given nationality, we plot $P(\widehat{x} = R \mid N)\,(1 - P(\widehat{x} = R \mid N))$ in Panel (c.2). Importantly, uncertainty follows an *inverted* U-shape: it peaks when the probability of red hair is 0.5, and converges to zero when hair color becomes certain given the nationality. The *opposing* patterns of these two effects (U-shaped and inverted U-shaped, respectively) lead to the non-monotonic, double hump-shaped effective overreaction. Initially, the news component dominates, but eventually the decline in uncertainty determines the convergence of the distorted probability to its objective counterpart. Finally, the sizes of the two humps in Panel (b.1) differ because the trough of the news component and the peak of the uncertainty effect do not coincide, implying

that when the probability of red hair increases, initially both the news and uncertainty components contribute to generating overreaction in the probability.

3. **Over/under confidence**. Overreaction to information also affects perceived uncertainty compared to objective uncertainty. Panel (b.2) reports the difference between perceived uncertainty, $P^\theta(\widehat{x} = R|N, W)(1 - P^\theta(\widehat{x} = R|N, W))$, and objective uncertainty $P(\widehat{x} = R|N)(1 - P(\widehat{x} = R|N))$. A negative value implies overconfidence: the agent is too certain about the hair color compared to the objective distribution.

The patterns of over/under confidence are driven by the interaction of the same news and uncertainty effects. First, there is no over- or underconfidence if there is no news, or the news completely removes uncertainty ($P(\widehat{x} = R|N) = \{0, 1\}$). Second, when $P(\widehat{x} = R|N)$ is smaller than the reference of 0.2, news and uncertainty effects work in the same direction, and only overconfidence is possible. Third, when instead $P(\widehat{x} = R|N)$ exceeds the reference value, the two effects can operate in opposite directions. For small increases in the probability of red hair relative to the reference, agent's overreaction to the new information can result in underconfidence. This occurs because the distorted probability shifts closer to 0.5—the point of maximal uncertainty. As the objective probability of red hair continues to rise, the influence of the news effect begins to dominate, leading to overconfidence. Eventually, as the news fully resolves uncertainty, the strength of overconfidence naturally declines, and the distortion effectively disappears.

3 Smooth Diagnostic Expectations

In this section, we show that the key principles uncovered for categorical distributions naturally find their counterparts under Smooth DE, which allows for continuous distributions and accumulation of information.

3.1 Smooth DE density: representativeness and information sets

We introduce the definition of representativeness and information sets with some generality. Consider the filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, P)$ where $(\mathcal{F}_t)_{t\geq 0}$ is an increasing family of sub-σ-algebras of $\mathcal{F}$, with $\mathcal{F}_t \subset \mathcal{F}_{t+1}$. The filtration $(\mathcal{F}_t)_{t\geq 0}$ represents the evolution of information over time. A stochastic process $(x_t)_{(t\geq 0)}$ on $(\Omega, \mathcal{F}, P)$ is adapted to the filtration $(\mathcal{F}_t)_{t\geq 0}$ if, for each $t \geq 0$, x_t is $\mathcal{F}_t$-measurable. Let the conditional density function of x_{t+h}, for some horizon $h \geq 1$, given the filtration $\mathcal{F}_t$, be denoted by $f(\widehat{x}_{t+h}|\mathcal{F}_t)$. This function describes the probability density of x_{t+h}, given all the information available up to time t.

In our approach, the time-series counterpart of equation (1) is obtained by defining the *reference group as a past available information set*. Mirroring the language of Tversky

and Kahneman (1975) in Section 2, in the time-series domain an event is representative for the current information set if the relative frequency of this event is higher conditional on the current information set than conditional on some past reference information set. Representativeness is then defined as follows:

Definition 1 ***(Representativeness and information sets)*** *Given the current information set $\mathcal{F}_t$ and a reference information set $\mathcal{F}_{t-J}$, formed $J \geq 1$ periods ago, the representativeness of a random event $\widehat{x}_{t+h}$ for some future horizon $h \geq 1$ is defined as*

$$rep\left(\widehat{x}_{t+h}|\mathcal{F}_t, \mathcal{F}_{t-J}\right) \equiv \frac{f\left(\widehat{x}_{t+h}|\mathcal{F}_t\right)}{f\left(\widehat{x}_{t+h}|\mathcal{F}_{t-J}\right)}. \tag{6}$$

Based on this definition of the representativeness of an event, we then follow the same construction of the distorted belief as introduced earlier in equation (2) in Section 2. In particular, we build the following distorted conditional density, which we refer to as *Smooth Diagnostic Expectations*, or in a more abbreviated form as Smooth DE.

Definition 2 ***(Smooth Diagnostic Expectations)****. In the time series domain, the conditional density $f^{\theta}\left(\widehat{x}_{t+h}|\mathcal{F}_t, \mathcal{F}_{t-J}\right)$, distorted by representativeness as defined in Definition 1, is constructed as follows*

$$f^{\theta}\left(\widehat{x}_{t+h}|\mathcal{F}_t, \mathcal{F}_{t-J}\right) = f\left(\widehat{x}_{t+h}|\mathcal{F}_t\right)\left[rep\left(\widehat{x}_{t+h}|\mathcal{F}_t, \mathcal{F}_{t-J}\right)\right]^{\theta}/Z \tag{7}$$

where Z is a constant of integration and the parameter $\theta \geq 0$ measures the distortion severity.

In the rest of this paper we focus on Normal densities, which allows for a closed-form distorted distribution. Under Normality, two moments, the conditional mean $\mathbb{E}_s(x_{t+h})$ and conditional variance $\mathbb{V}_s(x_{t+h})$, summarize the whole conditioning information:

$$f\left(x_{t+h}|\mathcal{F}_s\right) = \mathcal{N}\left(x_{t+h}; \mathbb{E}_s(x_{t+h}), \mathbb{V}_s(x_{t+h})\right), \quad \text{for} \quad s \leq t.$$

It follows that given Definition 1, the representativeness of the event $\widehat{x}_{t+h}$ is evaluated as

$$rep\left(\widehat{x}_{t+h}|\mathcal{F}_t, \mathcal{F}_{t-J}\right) = \frac{\mathcal{N}(\widehat{x}_{t+h}; \mathbb{E}_t(x_{t+h}), \mathbb{V}_t(x_{t+h}))}{\mathcal{N}(\widehat{x}_{t+h}; \mathbb{E}_{t-J}(x_{t+h}), \mathbb{V}_{t-J}(x_{t+h}))}, \tag{8}$$

capturing its relative frequency under the current information set $\mathcal{F}_t$, compared to some past information set $\mathcal{F}_{t-J}$, with $J \geq 1$. Then, the Smooth DE density in equation (7) has a closed-form solution given by Proposition 1 below.

Proposition 1 *Denote the ratio of variances for the current and reference groups as*

$$R_{t+h|t,t-J} \equiv \mathbb{V}_t(x_{t+h})/\mathbb{V}_{t-J}(x_{t+h}) \tag{9}$$

If $R_{t+h|t,t-J} < (1+\theta)/\theta$*, the Smooth DE density* $f^{\theta}(x_{t+h}|\mathcal{F}_t, \mathcal{F}_{t-J})$ *in equation (7) is also Normal. Its conditional mean is*

$$\mathbb{E}_t^{\theta}(x_{t+h}) = \mathbb{E}_t(x_{t+h}) + \widetilde{\theta}_{t+h|t,t-J}\left[\mathbb{E}_t(x_{t+h}) - \mathbb{E}_{t-J}(x_{t+h})\right] \tag{10}$$

where its effective overreaction to news $\widetilde{\theta}_{t+h|t,t-J}$ *is given by*

$$\widetilde{\theta}_{t+h|t,t-J} = \theta \frac{R_{t+h|t,t-J}}{1+\theta\left(1-R_{t+h|t,t-J}\right)} \tag{11}$$

The Smooth DE conditional variance is

$$\mathbb{V}_t^{\theta}(x_{t+h}) = \frac{\mathbb{V}_t(x_{t+h})}{1+\theta\left(1-R_{t+h|t,t-J}\right)} \tag{12}$$

Proof. See Appendix B.[4] ■

Comparison to standard DE. Our Smooth DE approach differs from the standard one in the literature, given by the Bordalo et al. (2018) (BGS) formulation, where the representativeness of $\widehat{x}_{t+h}$ is defined as the relative frequency

$$rep^{BGS}(\widehat{x}_{t+h}|\mathcal{F}_t, \mathcal{F}_{t-J}) \equiv \frac{\mathcal{N}(\widehat{x}_{t+h}; \mathbb{E}_t(x_{t+h}), \mathbb{V}_t(x_{t+h}))}{\mathcal{N}(\widehat{x}_{t+h}; \mathbb{E}_{t-J}(x_{t+h}), \mathbb{V}_t(x_{t+h}))}. \tag{13}$$

In equation (13), the reference density uses the mean conditional on the information set $\mathcal{F}_{t-J}$, $\mathbb{E}_{t-J}(x_{t+h})$, but the uncertainty conditional on the current information set $\mathcal{F}_t$, $\mathbb{V}_t(x_{t+h})$. Instead, under Smooth DE, the reference density is based entirely on the past information set $\mathcal{F}_{t-J}$, thus capturing the entire change in information between $t-J$ and t.

Bordalo et al. (2018) show how, when $\mathbb{V}_t(x_{t+h}) > 0$, the BGS formulation of representativeness in equation (13) implies that the DE density $f^{\theta}(\widehat{x}_{t+h}|\mathcal{F}_t, \mathcal{F}_{t-J})$ defined by equation

[4]The condition $R_{t+h|t,t-J} < (1+\theta)/\theta$ guarantees that the variance of the distorted Normal distribution is finite and positive. As $R_{t+h|t,t-J}$ approaches this limit, the variance of the Smooth DE distribution approaches infinity and the corresponding Normal distribution approaches a degenerate, flat distribution. In Appendix B we further discuss an approach that deals with this threshold condition by implementing an upper bound on the Smooth DE overreaction in the conditional mean. The approach guarantees that both the mean and variance distortions remain finite and non-decreasing as the ratio $R_{t+h|t,t-J}$ goes to infinity.

(7) has a Normal distribution with conditional mean

$$\mathbb{E}_t^{BGS,\theta}(x_{t+h}) = \mathbb{E}_t(x_{t+h}) + \theta\left[\mathbb{E}_t(x_{t+h}) - \mathbb{E}_{t-J}(x_{t+h})\right]$$

and conditional variance

$$\mathbb{V}_t^{BGS,\theta}(x_{t+h}) = \mathbb{V}_t(x_{t+h}).$$

Smooth DE and the standard BGS formulation of DE coincide in two cases. First, in the extreme case of no conditional uncertainty. Intuitively, when $\mathbb{V}_t(x_{t+h}) = 0$, the conditional likelihood of observing any other scenario for x_{t+h} than the one the agent is now fully informed on has become equal to zero. According to Proposition 1, Smooth DE nests that extreme, which amounts to $R_{t+h|t,t-J} = 0$, and thus effectively no distortion even if $\theta > 0$. In the BGS formulation that limit is imposed through a discontinuity at $\mathbb{V}_t(x_{t+h}) = 0$, where $\mathbb{E}_t^{BGS,\theta}(x_{t+h})$ is assumed to collapse to $\mathbb{E}_t(x_{t+h})$ (as in Bordalo et al. (2018)). Under Smooth DE the effective overreaction $\widetilde{\theta}_{t+h|t,t-J}$ *smoothly* goes to zero as current uncertainty goes to zero. Second, and more importantly, away from the $\mathbb{V}_t(x_{t+h}) = 0$ case, in the original BGS formulation the ratio $R_{t+h|t,t-J}$ is always 1 by assumption. Thus, DE coincides with Smooth DE if and only if the stochastic process is characterized by information sets $\mathcal{F}_t$ and $\mathcal{F}_{t-J}$ that happen to deliver $\mathbb{V}_t(x_{t+h}) = \mathbb{V}_{t-J}(x_{t+h}), \forall t$ and for any given J.

The fact that uncertainty influences the severity of the DE distortion under Smooth DE also leads to differences compared to the standard BGS formulation in terms of technical properties like the failure of the law of iterated expectations and the additivity under the Smooth DE operator. Appendix C provides details and proofs on these statistical features. These properties matter, for example for time-consistency of choices (see also the discussion in Bianchi et al. (2024) in the context of standard DE) when taking the Smooth DE to structural optimizing models, like in our business cycle application of Section 5.

3.2 Uncertainty, overreaction and confidence

In the following corollary we summarize three key properties of forecasts under Smooth DE. These properties mirror the qualitative ones for the static categorical example discussed in Section 2: 'overreaction', 'over/under confidence' and 'news and uncertainty', respectively. The closed-form solution of Normal densities makes these properties particularly transparent. This connection is useful, indicating that Smooth DE maintains similar insights when extending the concept of representativeness and overreaction to new information to dynamic models. Additionally, the tractability of the Normal distribution in the time series domain will allow us to connect these three properties to stylized survey facts in Section 3.3.1.

Corollary 1 *Assume the presence of residual uncertainty with respect to a future event:* $\mathbb{V}_t(x_{t+h}) > 0$. *Compared to the RE forecast* ($\theta = 0$)*, the conditional forecast under Smooth DE* ($\theta > 0$)*, characterized in Proposition 1, exhibits*

1. *overreaction of the conditional mean to new information, since* $\widetilde{\theta}_{t+h|t,t-J} > 0$.
2. *(i) overconfidence, i.e.* $\mathbb{V}_t^{\theta}(x_{t+h}) < \mathbb{V}_t(x_{t+h})$*, when* $R_{t+h|t,t-J} < 1$*, or (ii) underconfidence, i.e.* $\mathbb{V}_t^{\theta}(x_{t+h}) > \mathbb{V}_t(x_{t+h})$*, when* $R_{t+h|t,t-J} > 1$.
3. *an effective overreaction of the conditional mean to new information that is monotonically increasing in the ratio* $R_{t+h|t,t-J}$ *between current and past uncertainty*

$$\frac{\partial \widetilde{\theta}_{t+h|t,t-J}}{\partial R_{t+h|t,t-J}} > 0$$

To better understand Corollary 1, we consider a simple example and use a visual illustration. Suppose an agent is interested in forecasting the return of an asset at horizon h, denoted by x_{t+h}. Figure 2 presents four different cases illustrating how news received at time t affects the perceived density of future returns under DE and Smooth DE. The figure displays the reference, current, DE, and Smooth DE distributions. Appendix D provides additional details on these cases to illustrate how the news and uncertainty effects introduced in Section 2 determine the effective overreaction of the Smooth DE density.

Overconfidence. Consider first a case in which the new information implies only a reduction in uncertainty about future returns, while expected values remain unchanged: $\mathbb{E}_t(x_{t+h}) = \mathbb{E}_{t-J}(x_{t+h})$, $\mathbb{V}_t(x_{t+h}) < \mathbb{V}_{t-J}(x_{t+h})$, and thus $R_{t+h|t,t-J} < 1$. This case, shown in the top-left panel, serves as a natural benchmark because, for a wide range of stationary processes, a more recent information set typically leads to reduced uncertainty and more precise forecasts. This reduction implies that returns close to the mean—and the mean itself—become more likely, while tail events become less likely. Under the standard BGS implementation of DE, this change has no effect. In contrast, under Smooth DE, the agent revises her beliefs in light of the new information. She inflates the probability of the mean and other returns that have become more *representative*, while further downplaying the probability of tail events that have become less likely. The result is a distribution even narrower than the current one. Under Smooth DE, the new information leads to an *overreaction* in terms of the decline in uncertainty, resulting in a novel implication: overconfidence.

Underconfidence. Now consider the opposite case, in which uncertainty increases in response to recent news, despite the forecast horizon becoming shorter. This case is shown in the bottom-left panel. The current distribution now has a *larger* variance than the reference

Figure 2: Smooth Diagnostic Expectations and standard Diagnostic Expectations densities

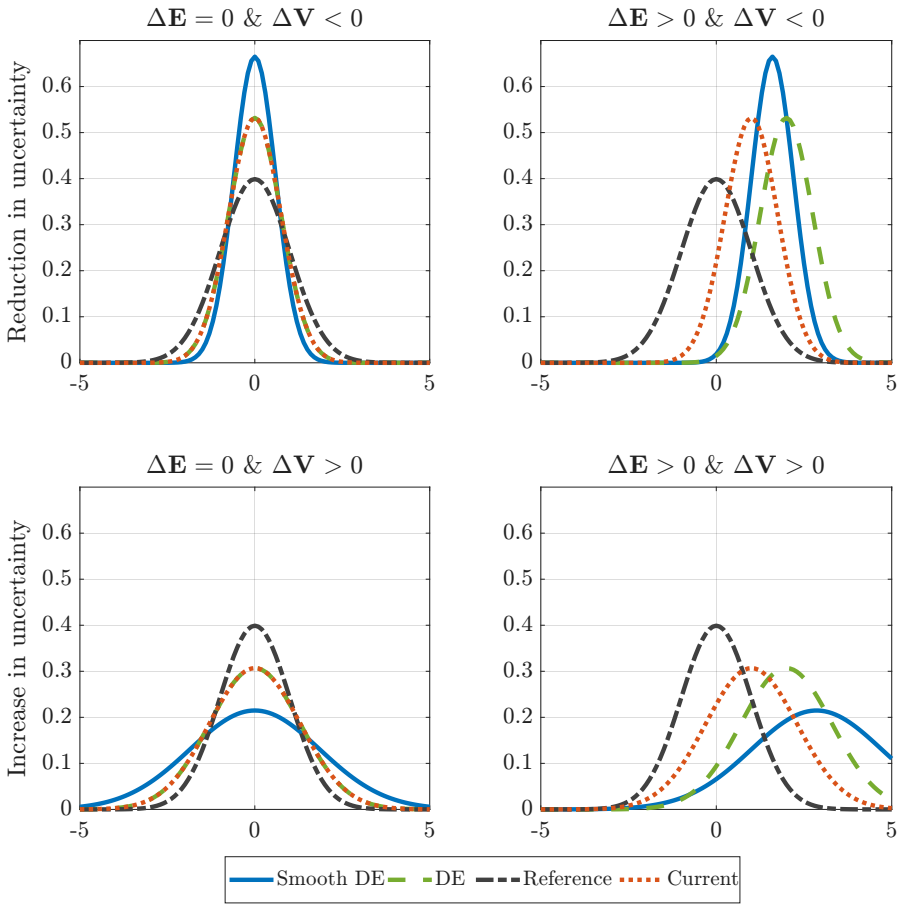

Notes: The figure uses the following mean and variance values of the current and reference distributions: $\mathbb{E}_{t-J}(x_{t+h}) = [0, 0, 0, 0]$, $\mathbb{E}_t(x_{t+h}) = [0, 0, 1, 1]$, $\mathbb{V}_{t-J}(x_{t+h}) = [1, 1, 1, 1]$, and $\mathbb{V}_t(x_{t+h}) = [.75, 1.3, .75, 1.3]$.

distribution, i.e., $\mathbb{V}_t(x_{t+h}) > \mathbb{V}_{t-J}(x_{t+h})$, and thus $R_{t+h|t,t-J} > 1$. Such a situation could arise, for example, in response to a positive uncertainty shock. Tail events become more representative under the revised density and receive magnified weight under Smooth DE. As a manifestation of overreaction to the new information, probability mass shifts from the center to the tails, and the agent becomes *underconfident* about future returns. In contrast, even in this case, the DE density coincides with the current density.

Taken together, Smooth DE exhibits a novel form of overreaction to changes in uncertainty. Even without a shift in the mean, new information is revealed and the probability distribution adjusts when the variance changes, as tail events become more or less likely. Thus, overconfidence and underconfidence can be interpreted as manifestations of overreaction, once we recognize that the entire density reflects agents' beliefs. Nevertheless, in what follows, we adopt the convention commonly used in the literature: we use *overreaction* to

refer specifically to a shift in the expected value of a variable, and *overconfidence* to describe a situation in which subjective uncertainty is lower than objective uncertainty.

Changes in uncertainty affect the overreaction of the conditional mean. The two panels in the right column of Figure 2 consider cases in which the new information implies a revision of both the mean and the variance of returns. Since changes in information sets typically involve actual changes in both conditional moments, this is a particularly novel and important aspect of Smooth DE.

In the top-right panel, the current distribution shifts to the right relative to the reference density, and uncertainty is reduced. The DE density also shifts to the right, but with no change in shape or uncertainty relative to the current distribution. Under Smooth DE, however, the agent recognizes that, despite the increase in the mean, the new information has made tail events to the right less representative. Thus, for a given θ, the Smooth DE density does shift to the right, but by a smaller amount, and becomes visibly narrower, reflecting the reduction in uncertainty. As the current conditional uncertainty continuously decreases and converges to zero, the shift in the distorted conditional mean also smoothly converges to zero. In contrast, under standard DE, there would be no effect of the change in uncertainty on the overreaction in mean, unless we reach the extreme case of $\mathbb{V}_t(x_{t+h}) = 0$ where the distortion discretely disappears.

Finally, new information may also bring a shift in the mean, but now with *increased* uncertainty. The bottom-right panel illustrates this case. The shift in the mean is exactly the same as in the top-right panel, but now the variance is larger rather than smaller. The agent's overreaction in terms of revisions to the conditional mean is now *stronger* than under standard DE because tail events have become more likely under the current distribution. This leads to an even more significant shift of probability mass to the right. As in the bottom-left panel, we observe underconfidence.

3.3 Two baseline stochastic environments

Proposition 1 presents the Smooth DE density, which takes as input the objective conditional moments characterizing a given stochastic environment and distorts them. In this section, we apply Smooth DE to two commonly used baseline environments. The first is a standard AR(1) process, which we analyze analytically to better understand the properties of Smooth DE, distinguish them from the standard DE, and connect them to stylized survey facts documented in the literature. The second environment introduces noisy signals about an otherwise AR(1)-driven state. This case allows us to derive a version of the standard Kalman Filter under the Smooth DE density, an approach that is particularly relevant for applied

time-series work. We also use this framework to link our findings to the new survey evidence presented in Section 4 and to the business cycle application discussed in Section 5.

3.3.1 AR(1) process

Consider a standard AR(1) process

$$x_t = \rho x_{t-1} + u_t, \qquad u_t \sim N(0, \sigma_u^2) \tag{14}$$

with $\rho < 1$ and $\sigma_u^2 > 0$. It follows that the objective conditional moments are given by $\mathbb{E}_{t-J}(x_{t+h}) = \rho^{h+J} x_{t-J}$ and $\mathbb{V}_{t-J}(x_{t+h}) = \sigma_u^2(1-\rho^{2(h+J)})/(1-\rho^2)$, for $J \geq 0$. It also follows that the key ratio $R_{t+h|t,t-J}$ of conditional variances in equation (9) is given by

$$R_{t+h|t,t-J} = \frac{1-\rho^{2h}}{1-\rho^{2(h+J)}} \tag{15}$$

By Proposition 1, it follows that for a horizon $h \geq 1$ and a reference distribution with $J \geq 1$, the conditional mean for the Smooth DE density $f^{\theta}(\widehat{x}_{t+h})$ for the AR(1) process is

$$\mathbb{E}_t^{\theta}(x_{t+h}) = \rho^h x_t + \widetilde{\theta}_{t+h|t,t-J}\left(\rho^h x_t - \rho^{h+J} x_{t-J}\right), \tag{16}$$

where, according to equation (11), the effective overreaction $\widetilde{\theta}_{t+h|t,t-J}$ increases with $R_{t+h|t,t-J}$.

Corollary 1 emphasizes the role of the ratio of conditional variances in driving the effective distortion. By analyzing the properties of $R_{t+h|t,t-J}$ in equation (15), we use the specific AR(1) example to further establish analytically novel properties compared to standard DE, and then connect them to existing survey evidence.

Corollary 2 *The ratio $R_{t+h|t,t-J} < 1$ for any $h \geq 1$ and $J \geq 1$. The ratio is (i) increasing in the horizon h and (ii) decreasing in the autoregressive coefficient ρ and memory lag J.*

Figure 3 illustrates these properties considering the effects of a unit shock for different ρ and J values. For each horizon h, we report the RE forecast, the effective DE distortion $\widetilde{\theta}_{t+h|t,t-J}$, the difference between the Smooth DE and the RE forecast, $\rho^h \widetilde{\theta}_{t+h|t,t-J}\epsilon_t$, and the Smooth DE forecast. We connect these properties to existing survey evidence as follows.

Overreaction *and* overconfidence. Intuitively, in stationary, homoskedastic environments, events closer into the future are typically easier to predict than events far into the future. In this AR(1) example, we see that effect as new information reduces conditional variances, leading to $R_{t+h|t,t-J} < 1$ in Corollary 2. By Corollary 1 the Smooth DE density is then characterized by subjective overconfidence, with $\mathbb{V}_t^{\theta}(x_{t+h}) < \mathbb{V}_t(x_{t+h})$. This

Figure 3: Response of expectations to a unit shock in an AR(1) process

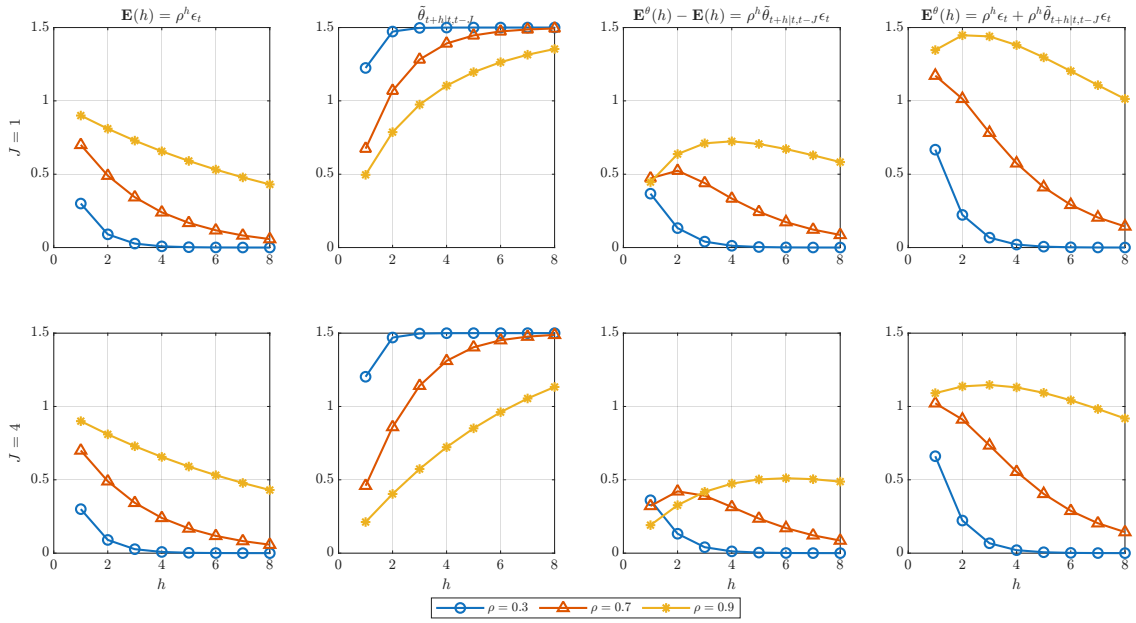

Notes: We plot the effects of a unit innovation to an AR(1) process for different values of ρ and J assuming $\theta = 1.5$. For each horizon h, the four columns report the RE forecast, the effective DE distortion $\widetilde{\theta}_{t+h|t,t-J}$, the difference between the Smooth DE and the RE forecast $\rho^h \widetilde{\theta}_{t+h|t,t-J}\epsilon_t$, and the Smooth DE forecast.

implication is consistent with stylized survey data. In fact, a recent literature studying the properties of survey responses, including Barrero (2022), Born et al. (2025), and the reviews in Altig et al. (2022) and Born et al. (2022), documents that firms not only *overreact* to news, but that they are also typically *overconfident* in their subjective forecasts.

Stronger overreaction for longer forecast horizons. For a given lag J in the reference distribution, as the forecast horizon h increases, the effective horizon of the current RE forecast (h), and the effective horizon of the reference RE forecast $(h + J)$ become increasingly similar. As a result, the levels of uncertainty associated with the two forecasts also become increasingly similar because, in relative terms, the current information set is less and less informative for the variable that the agent is trying to predict. Formally, in Corollary 2 the ratio $R_{t+h|t,t-J}$ increases towards one as the horizon increases toward infinity. Given that under Smooth DE overreaction is increasing in the level of relative uncertainty, the longer the horizon h, the stronger the overreaction to a given revision of the RE forecast. For the same reason, the effective overreaction decreases with the memory lag J. A longer memory lag J implies a decline in relative uncertainty, as the current information significantly reduces conditional uncertainty compared to the reference distribution.

A strand of literature using survey data argues that overreaction increases with the fore-

cast horizon, consistent with this implication of Smooth DE. Bordalo et al. (2019), Bordalo et al. (2020), d'Arienzo (2020), Bordalo et al. (2023), point to stronger overreaction for equity analysts' forecasts of long-term earning growth and professional forecasters' forecasts of long-term interest rates. Halperin and Mazlish (2025) document across a wide set of countries horizon-increasing overreaction of consensus forecasts for output, consumption, investment, and inflation. Bonaglia et al. (2025) focus on similar patterns in inflation and interest rate forecasts. Augenblick et al. (2025) use field data from betting and financial markets to argue that compared to the Bayesian forecast there is relatively stronger overreaction to signals with a longer time-to-resolution, conceptually similar to longer forecast horizons.

Stronger overreaction for less persistent processes. A lower autoregressive parameter ρ implies less predictability for future events and more similar levels of uncertainty between the current and reference distributions. Thus, while the last term in parentheses on the right-hand side of equation (16) gets smaller as ρ declines, a less persistent process also results in more overreaction *for a given revision of the objective conditional mean.* This implication of Smooth DE is qualitatively consistent with the evidence presented in Afrouzi et al. (2023), with the caveat that the environments considered in the experiments of that paper are more complex than those implied by the simple AR(1) example presented here.

3.3.2 Signal extraction under Smooth DE

In this section we extend our Smooth DE framework to the case where agents have noisy information about the current state. Formally, we start with the same probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, P)$ introduced in Section 3.1, but now we allow for noisy information. In particular, let the information available up to time t be represented by another filtration $(\mathcal{G}_t)_{t\geq 0}$, generated by an observed process s_t of signals. Then let $f(\widehat{x}_{t+h}|\mathcal{G}_t)$ be the conditional density function of x_{t+h} based on the information available in $\mathcal{G}_t$ from the imperfect observations.

Like in Definitions 1 and 2, the representativeness of an event is

$$rep\left(\widehat{x}_{t+h}|\mathcal{G}_t, \mathcal{G}_{t-J}\right) \equiv \frac{f\left(\widehat{x}_{t+h}|\mathcal{G}_t\right)}{f\left(\widehat{x}_{t+h}|\mathcal{G}_{t-J}\right)} \tag{17}$$

while the conditional density distorted by representativeness is

$$f^{\theta}\left(\widehat{x}_{t+h}|\mathcal{G}_t, \mathcal{G}_{t-J}\right) = f\left(\widehat{x}_{t+h}|\mathcal{G}_t\right)\left[rep\left(\widehat{x}_{t+h}|\mathcal{G}_t, \mathcal{G}_{t-J}\right)\right]^{\theta} Z^{-1} \tag{18}$$

where Z is a constant of integration and the parameter $\theta \geq 0$. Horizon $h = 0$ corresponds to a nowcast, i.e. the conditional density of the current unobserved x_t.

To derive closed-form solutions, we further assume a standard Gaussian case of noisy

information, maintain $J = 1$ and focus on the current conditional estimate at $h = 0$. In particular, let the unobserved x_t follow the same AR(1) process as in equation (14), and now add the observation equation with the noisy signal s_t:

$$s_t = x_t + \varepsilon_t, \qquad \varepsilon_t \sim N(0, \sigma_\varepsilon^2) \tag{19}$$

The Smooth DE density distorts the rational, or Bayesian, period t density $f\left(\widehat{x}_t|\mathcal{G}_t\right)$. The Kalman Filter delivers a recursive law of motion for the Bayesian conditional mean $\mathbb{E}_t(x_t)$

$$\mathbb{E}_t(x_t) = \rho\mathbb{E}_{t-1}(x_{t-1}) + K_t\left(s_t - \rho\mathbb{E}_{t-1}(x_{t-1})\right), \tag{20}$$

where the response to the news $(s_t - \rho\mathbb{E}_{t-1}(x_{t-1}))$ depends on the Kalman gain

$$K_t = \frac{\rho^2\mathbb{V}_{t-1}(x_{t-1}) + \sigma_u^2}{\rho^2\mathbb{V}_{t-1}(x_{t-1}) + \sigma_u^2 + \sigma_\varepsilon^2}, \tag{21}$$

and a recursive update of the posterior uncertainty $\mathbb{V}_t(x_t)$ given by $\mathbb{V}_t(x_t) = K_t\sigma_\varepsilon^2$. Key for Smooth DE, the ratio of posterior and prior uncertainty over x_t is then simply

$$R_{t|t,t-1} = \frac{\mathbb{V}_t(x_t)}{\rho^2\mathbb{V}_{t-1}(x_{t-1}) + \sigma_u^2} \tag{22}$$

The Smooth DE density then follows from Proposition 1, taking as inputs these Bayesian conditional moments.[5] In particular, the distorted conditional mean follows from equation (10) and that $\mathbb{E}_{t-1}(x_t) = \rho\mathbb{E}_{t-1}(x_{t-1})$, as

$$\mathbb{E}_t^\theta\left(x_t\right) = \mathbb{E}_t(x_t) + \widetilde{\theta}_{t|t,t-1}\left(\mathbb{E}_t(x_t) - \rho\mathbb{E}_{t-1}(x_{t-1})\right)$$

By using the recursion of the Bayesian nowcast in equation (20) the distorted mean is thus

$$\mathbb{E}_t^\theta\left(x_t\right) = \mathbb{E}_t(x_t) + \widetilde{\theta}_{t|t,t-1}K_t\left(s_t - \rho\mathbb{E}_{t-1}(x_{t-1})\right), \tag{23}$$

where $\widetilde{\theta}_{t|t,t-1}$ is again given by equation (11), increasing in the ratio $R_{t|t,t-1}$ in (22). The Smooth DE agents also update their belief about x_t with news in proportion to the Kalman gain K_t but they overreact relative to the Bayesian benchmark by $\widetilde{\theta}_{t|t,t-1}$.

[5]Appendix F offers more details. Our derivation of the Kalman filter used by agents subject to Smooth DE connects and extends the diagnostic Kalman Filter derived within the standard BGS formulation in earlier work like Bordalo et al. (2019) and Bordalo et al. (2020).

4 New survey evidence: Overreaction and Uncertainty

In the previous section we have shown that Smooth DE offers a microfoundation for known survey patterns like overreaction and overconfidence. We now turn to empirical evidence from survey forecasts to assess the relevance of this mechanism. We uncover a novel stylized fact that strongly supports the theory: individual forecast overreaction significantly increases when uncertainty is high (Corollary 1).

Data. We use the quarterly individual-level SPF forecast data (1968:Q4–2023:Q2). The surveys are collected in the second month of each quarter and report forecasts of the current and the next four quarters' outcomes of the variables in question. We study forecasts of 15 macro variables, including GDP, consumption, and investment, CPI, unemployment rate, and short- and long-term interest rates. We follow Bordalo et al. (2020) exactly in defining the variables and thus defer to their Online Appendix B for further details.

We follow Bordalo et al. (2020) and focus on a one-year-ahead forecast horizon. Since most SPF forecasts are in levels, for variables such as GDP and prices, we convert them into implied growth rates by taking growth rates from quarter $t-1$ levels to quarter $t+3$ forecasts. For variables such as the unemployment rate and interest rates, we study the $t+3$-quarter-ahead levels. We follow Bordalo et al. (2020) and remove outliers by dropping forecasts that are 5 interquartile ranges away from the median forecast for each quarter. We then drop forecasters who appeared in the survey less than 10 times. We perform the cleaning procedure before computing forecast errors, revisions, and cross-sectional standard deviations. The data cleaning procedure results in 23 independent forecaster observations on average per quarter for each series. Next, consider the data on actual outcomes. Macroeconomic variables are often revised after the initial release. To capture the forecasters' information set as closely as possible, we use the initial releases from the Philadelphia Fed's Real-Time Data Set for Macroeconomists. Since financial data are not revised, for the interest rates we use historical data from the St. Louis Fed's FRED.

Empirical strategy. We follow Bordalo et al. (2020), who build on Coibion and Gorodnichenko (2015), and study the correlation between forecast revisions and subsequent forecast errors. Our innovation is to examine how this correlation depends on the level of uncertainty, providing a test for Smooth DE. We denote the h-quarters-ahead forecasts at time t of a variable x by a forecaster i as $x_{i,t+h|t}$. The consensus forecast $x_{t+h|t}$ is given by the mean forecast across forecasters: $x_{t+h|t} = \frac{1}{I}\sum_{i=1}^{I} x_{i,t+h|t}$, where I is the total number of forecasters.

Consensus regression. First, consider the Coibion and Gorodnichenko (2015) regression of consensus forecast errors, $x_{t+h} - x_{t+h|t}$, on consensus forecast revisions $x_{t+h|t} - x_{t+h|t-1}$:

$$x_{t+h} - x_{t+h|t} = \alpha + \beta_C \left(x_{t+h|t} - x_{t+h|t-1}\right) + \epsilon_{t,t+h}. \tag{24}$$

Under full-information-rational-expectations (FIRE), forecast errors are unpredictable so $\beta_C = 0$. Negative predictability $\beta_C < 0$ indicates consensus overreaction. Intuitively, when a forecast revision is positive, the forecast error is on average negative and hence the forecast is too optimistic. Conversely, $\beta_C > 0$ indicates consensus underreaction.

Individual-level regression. The consensus regression does not directly test rationality, as aggregation effects can induce predictability even under RE (Coibion and Gorodnichenko (2015)). Thus, Bordalo et al. (2020) extend (24) to examine the predictability of forecast errors at the individual level. They pool observations across all forecasters and estimate

$$x_{t+h} - x_{i,t+h|t} = \alpha_i + \beta_I \left(x_{i,t+h|t} - x_{i,t+h|t-1} \right) + \epsilon_{i,t,t+h}. \tag{25}$$

Rationality implies that forecast errors are unpredictable at the individual level, $\beta_I = 0$, even with information frictions. In contrast, if individual forecasts overreact relative to RE, we have $\beta_I < 0$, and if individual forecasts underreact, we have $\beta_I > 0$. Bordalo et al. (2020) find $\beta_I < 0$ for many macroeconomic variables, consistent with overreaction.

State-dependent individual-level regression. We allow the forecast error predictability in (25) to be state-dependent in order to examine whether and how the degree of overreaction depends on the level of uncertainty:

$$x_{t+h} - x_{i,t+h|t} = \alpha_i + \left(\gamma_I + \mathbf{1}\{\sigma_t = \sigma^H\} \times \delta_I \right) \left(x_{i,t+h|t} - x_{i,t+h|t-1} \right) + \epsilon_{i,t,t+h}, \tag{26}$$

where $\mathbf{1}\{\sigma_t = \sigma^H\}$ is a dummy variable that takes the value of one when uncertainty is high ($\sigma_t = \sigma^H$) and zero otherwise. We focus on the state-dependence in the individual-level regressions (25) rather than in the consensus regressions (24) for two reasons. First, we are interested in the degree of departures from RE, whereas the consensus regression does not directly test for rationality. Second, the individual-level regression increases the number of observations in each quarter relative to the consensus regression, thereby enhancing statistical power and making it more suitable for studying the implications of infrequent episodes of sharp increases in uncertainty as documented in Bloom (2009).

The overall predictability of forecast errors is given by γ_I in normal times and by $(\gamma_I + \delta_I)$ when uncertainty is high. Thus δ_I captures the change in predictability during high uncertainty periods. This specification allows a direct test of Smooth DE against standard DE. Standard DE, featuring constant overreaction, predicts $\delta_I = 0$. In contrast, Smooth DE predicts that overreaction intensifies with uncertainty (Corollary 1), implying $\delta_I < 0$.

We measure uncertainty using two approaches. In the first approach, we use the cross-sectional standard deviation of SPF forecasts of the variable of interest. The cross-sectional dispersion of survey forecasts has been used as a measure of uncertainty in several studies,

Figure 4: Number of series with cross-sectional dispersion significantly exceeding its mean

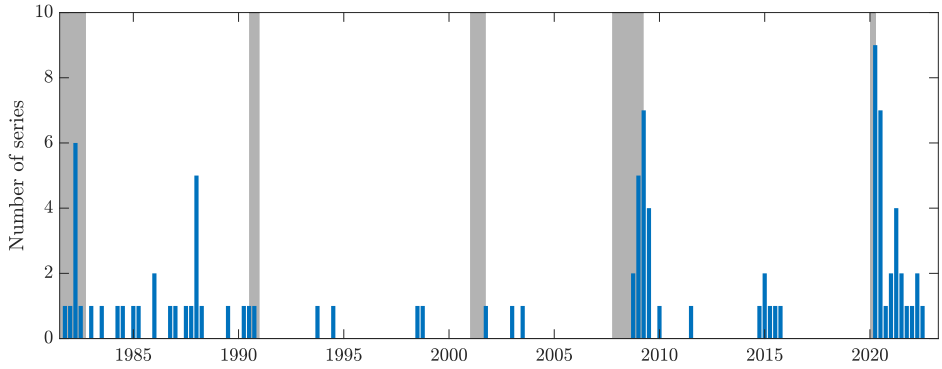

Notes: The figure plots, from 1981:Q3 to 2023:Q2, the number of series (out of 14 series) with cross-sectional standard deviations exceeding their time-series means by more than 1.65 standard deviations. The shaded areas are NBER recession dates. We start the graph from 1981:Q3 instead of 1968:Q4 since several SPF forecasts are only available from 1981:Q3. We exclude the Ten-year Treasury rate (TN10Y) from the graph since its SPF forecast is available only from 1992:Q1. As a result, we are left with 14 instead of 15 series.

such as Bachmann et al. (2013) and Ilut and Schneider (2014). This choice is internally consistent under Smooth DE, because higher uncertainty translates into greater cross-sectional dispersion of forecasts due to state-dependent overreaction, a channel absent in RE or standard DE, as we show below. We remove low-frequency movements from the raw time series of the cross-sectional dispersion using the Hodrick-Prescott (HP) filter ($\lambda = 1600$). We then identify high uncertainty periods as the ones during which the detrended disagreement is in the top 5%, corresponding to values 1.65 standard deviations above the mean.

At each point in time, Figure 4 shows how many of the 14 series we examine exhibit high uncertainty. There are four notable periods when this number rises significantly: the 1981 recession, the immediate aftermath of Black Monday (1988:Q1), the 2007–2009 Great Recession, and the COVID-19 recession and its aftermath. The figure illustrates a clear tendency for uncertainty to increase during recessions. This countercyclical pattern highlights the macroeconomic relevance of our mechanism, as periods of elevated uncertainty and subdued economic activity are precisely when Smooth DE predicts stronger overreaction.

In the second approach, we use the 3-month-ahead macroeconomic uncertainty index developed by Jurado et al. (2015), which is based on a large set of macroeconomic time series and is available from the St. Louis Fed (FRED). We apply the same procedure described above to identify periods of high uncertainty. In this case, however, the labeling of high uncertainty periods is uniform across all series, since it is derived from a single aggregate uncertainty index rather than from variable-specific cross-sectional dispersion.

Results. Figure 5 plots the estimation results for the individual-level regressions (25)

Figure 5: Forecast error on forecast revision regression results

(a) Cross-sectional standard deviation

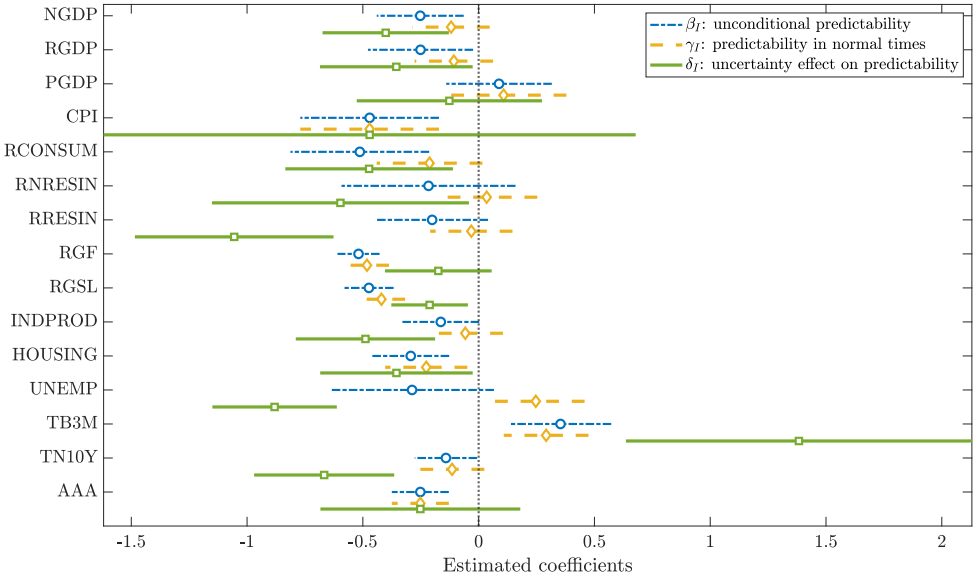

(b) Jurado et al. (2015) uncertainty index

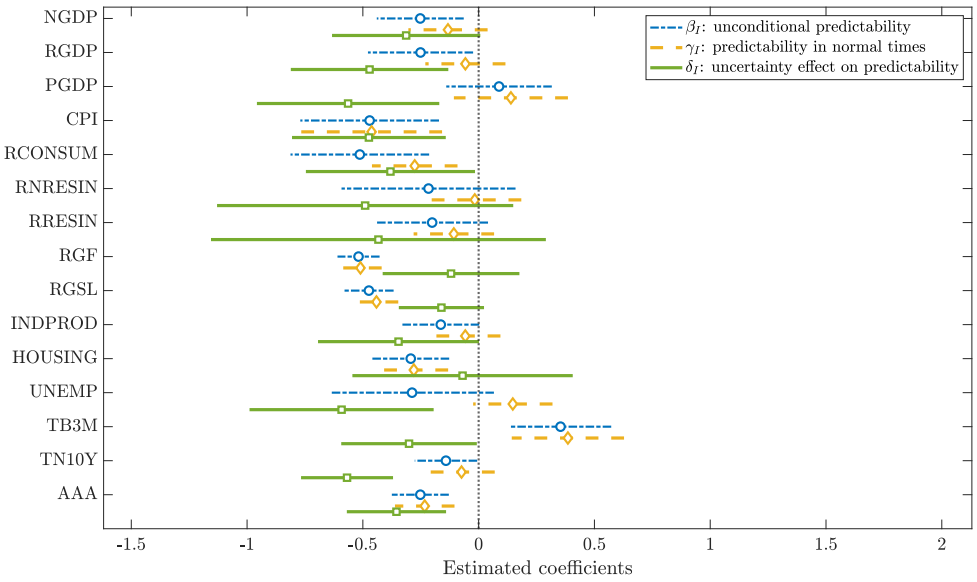

Notes: The figures report the estimates from the individual-level forecast error on forecast revision regressions. The blue circles represent β_I in the regression (25) and the yellow diamonds and green lines represent γ_I and δ_I in the regression with uncertainty dummy (26). In Panel (a) and (b), we measure uncertainty using the cross-sectional standard deviation of forecasts and Jurado et al. (2015) 3-month-ahead macroeconomic uncertainty index, respectively. The bars represent 90% confidence intervals where standard errors are Newey and West (1994) with automatic lag selection and clustered by forecaster and time.

and (26), and Table 1 reports the point estimates and standard errors for (24), (25), and (26). Standard errors are Newey and West (1994) with automatic lag selection and clustered by forecaster and time for individual-level regressions (25) and (26). We control for individual forecaster fixed effects in (25) and (26). The blue circles in Figure 5 and the second column in Table 1 report the forecast error predictability from the regression (25) without state-dependence. The point estimates for the unconditional predictability β_I are negative for 13 out of 15 series, with 8 of them significant at the 5 percent level. The results indicate that survey expectations exhibit individual-level overreaction, consistent with Bordalo et al. (2020)'s finding using earlier samples. The first column of Table 1 reports the estimates from the consensus regression (24). The point estimates of β_C are positive for 11 out of 15 series, but are significantly positive at the 5 percent level for only 2 of them.

The yellow diamonds and green squares in Panel (a) of Figure 5 and the third and fourth columns in Table 1 report the estimates from the uncertainty dummy regression (26), where uncertainty is measured using the cross-sectional standard deviation of SPF forecasts. The estimates strongly indicate state-dependent overreaction in survey forecasts. The estimated δ_I is typically negative, which implies that forecasters overreact more when uncertainty is high. Indeed, the point estimates of δ_I are negative for all series except for the 3-month T-bill rate. The estimated δ_I is significantly negative for 7 out of 15 series at the 5 percent level and for 3 additional series at the 10 percent level. Those variables include NIPA variables such as nominal and real GDP, personal consumption, real nonresidential and residential investment, as well as industrial production, unemployment rate, and the 10-year rate. Panel (b) of Figure 5 and the fifth and sixth columns in Table 1 report the case where we use the Jurado et al. (2015) 3-month-ahead macroeconomic uncertainty index. Similar to the first specification, the estimates support the hypothesis of stronger overreaction during periods of high uncertainty: the estimated δ_I is significantly negative for 6 out of 15 series at the 5 percent level and for 2 additional series at the 10 percent level. Interestingly, in contrast to the first specification, δ_I for the 3-month T-bill rate is negative at 10 percent level.

Compared to Bordalo et al. (2020), our regression results display stronger unconditional overreaction both at the consensus and individual levels. In Bordalo et al. (2020), β_C (consensus forecast error predictability) is significantly *positive* at the 5 percent level for 7 out of 15 SPF series (2 out of 15 series in our sample) and β_I (individual-level unconditional forecast error predictability) is significantly *negative* at the 5 percent level for 6 out of 15 SPF series (8 out of 15 series in our sample). We find stronger evidence in favor of overreaction because our sample includes the COVID-19 recession and its aftermath, a period of high uncertainty (Figure 4). Furthermore, comparing the point estimates of β_I with γ_I (forecast error predictability in normal times), for many variables, including nominal and

real GDP, consumption, and industrial production, γ_I is less negative than β_I. In addition, for these variables γ_I is estimated less precisely than β_I. This implies that the estimated unconditional overreaction β_I for those series is at least partially driven by strong overreaction during periods of high uncertainty.

In Appendix E, we check the robustness of these results by conducting aggressive treatments of outliers as advocated in Angeletos et al. (2021). Our main empirical finding is unaffected: survey forecasts exhibit stronger overreaction during periods of high uncertainty.

Summarizing, SPF forecasts exhibit (1) non-negative forecast error predictability at the consensus level, (2) overreaction at the individual level, and (3) stronger overreaction during periods of high uncertainty. We next show that Smooth DE can explain all three properties, while standard DE can only explain (1) and (2).

Forecasting framework. To derive analytical results on the key regression coefficients, we build on the filtering problem of Section 3.3.2. We implement a minimal extension of that standard framework to study the effect of time-varying uncertainty on forecast overreaction, a central feature also present in our subsequent RBC analysis (Section 5).

In particular, we let each agent i forecast the h-period-ahead future values of y_t, which is the sum of the hidden persistent component x_t *and* an i.i.d. temporary component, η_t:

$$y_t = x_t + \eta_t, \quad \eta_t \sim N(0, \sigma^2_{\eta,t}),$$

where the volatility $\sigma_{\eta,t}$ of the temporary component is *time-varying* and takes two values $\sigma_{\eta,t} \in \left\{\sigma^L, \sigma^H\right\}$, with $\sigma^L < \sigma^H$. All forecasters know the current realizations of $\sigma_{\eta,t}$. Like in Section 3.3.2, a forecaster i receives a private noisy signal $s_{i,t}$ on x_t as in equation (19), where the state x_t follows the AR(1) process in equation (14). We thus note that we deliberately continue to assume that the variances σ^2_u and σ^2_ε of shocks to the hidden state x_t and noise, respectively, are constant. Consequently, any time-variation in effective overreaction stems from the Smooth DE mechanism and not from changes in the noise-to-signal ratio.

Proposition 1 and equation (23) imply the Smooth DE forecast of y_{t+h} at period t

$$\mathbb{E}^\theta_{i,t}\left[y_{t+h}\right] = \rho^h \left[\mathbb{E}_{i,t}(x_t) + \widetilde{\theta}_{t+h|t,t-1} K \left(s_{i,t} - \rho \mathbb{E}_{i,t-1}(x_{t-1})\right)\right], \tag{27}$$

where the key ratio of conditional uncertainties is

$$R_{t+h|t,t-1} = \frac{\mathbb{V}_t\left(y_{t+h}\right)}{\mathbb{V}_{t-1}\left(y_{t+h}\right)} = \frac{\mathbb{V}_t\left(x_{t+h}\right) + \mathbb{V}_t\left(\eta_{t+h}\right)}{\mathbb{V}_{t-1}\left(x_{t+h}\right) + \mathbb{V}_{t-1}\left(\eta_{t+h}\right)}. \tag{28}$$

The posterior variances $\mathbb{V}_t\left(x_{t+h}\right)$ and $\mathbb{V}_{t-1}\left(x_{t+h}\right)$ of the persistent component are time invariant. Instead, stochastic volatility in the temporary component, η_t, causes the ratio

Table 1: Forecast error on forecast revision regression coefficients

	(1)	(2)	(3)	(4)	(5)	(6)
			Cross-sectional std		JLN uncertainty	
	β_C	β_I	γ_I	δ_I	γ_I	δ_I
Nominal GDP	0.25	-0.25^{**}	-0.12	-0.40^{**}	-0.13	-0.31
(NGDP)	(0.38)	(0.11)	(0.10)	(0.17)	(0.10)	(0.19)
Real GDP	0.11	-0.25^{*}	-0.11	-0.36^{*}	-0.06	-0.47^{**}
(RGDP)	(0.33)	(0.14)	(0.10)	(0.20)	(0.10)	(0.21)
GDP price index	1.52^{***}	0.09	0.11	-0.13	0.14	-0.56^{**}
(PGDP)	(0.44)	(0.14)	(0.16)	(0.24)	(0.15)	(0.24)
CPI	0.16	-0.47^{***}	-0.47^{***}	-0.47	-0.46^{**}	-0.47^{**}
(CPI)	(0.23)	(0.18)	(0.18)	(0.70)	(0.18)	(0.20)
Real consumption	-0.29	-0.51^{***}	-0.21	-0.47^{**}	-0.28^{**}	-0.38^{*}
(RCONSUM)	(0.29)	(0.18)	(0.14)	(0.22)	(0.11)	(0.22)
Real nonresidential investment	0.36	-0.22	0.03	-0.60^{*}	-0.02	-0.49
(RNRESIN)	(0.47)	(0.23)	(0.13)	(0.34)	(0.12)	(0.39)
Real residential investment	0.63	-0.20	-0.03	-1.06^{***}	-0.11	-0.43
(RRESIN)	(0.46)	(0.15)	(0.11)	(0.26)	(0.11)	(0.44)
Real federal government cons	-0.39^{*}	-0.52^{***}	-0.48^{***}	-0.17	-0.51^{***}	-0.12
(RGF)	(0.22)	(0.06)	(0.06)	(0.14)	(0.06)	(0.18)
Real state & local govmt cons	0.15	-0.47^{***}	-0.42^{***}	-0.21^{**}	-0.44^{***}	-0.16
(RGSL)	(0.31)	(0.06)	(0.06)	(0.10)	(0.06)	(0.11)
Industrial production	0.55^{*}	-0.16^{*}	-0.06	-0.49^{***}	-0.06	-0.35
(INDPROD)	(0.32)	(0.10)	(0.10)	(0.18)	(0.09)	(0.21)
Housing start	0.19	-0.29^{***}	-0.23^{**}	-0.35^{*}	-0.28^{***}	-0.07
(HOUSING)	(0.33)	(0.10)	(0.11)	(0.20)	(0.09)	(0.29)
Unemployment rate	-0.08	-0.29	0.25^{**}	-0.88^{***}	0.15	-0.59^{**}
(UNEMP)	(0.33)	(0.21)	(0.13)	(0.16)	(0.10)	(0.24)
Three-month Treasury rate	0.78^{***}	0.35^{***}	0.29^{***}	1.38^{***}	0.39^{***}	-0.30^{*}
(TB3M)	(0.23)	(0.13)	(0.11)	(0.45)	(0.15)	(0.18)
Ten-year Treasury rate	0.14	-0.14^{*}	-0.11	-0.67^{***}	-0.07	-0.57^{***}
(TN10Y)	(0.16)	(0.08)	(0.08)	(0.18)	(0.09)	(0.12)
AAA corporate bond rate	-0.03	-0.25^{***}	-0.25^{***}	-0.25	-0.23^{***}	-0.35^{***}
(AAA)	(0.22)	(0.07)	(0.07)	(0.26)	(0.08)	(0.13)

Notes: The table reports estimates from the consensus and individual-level forecast error on forecast revision regressions. The first column reports the coefficient β_C in the consensus regression (24) and the second column the coefficient β_I in the individual-level regression (25). The third to sixth columns report the coefficients γ_I and δ_I in the individual-level regression with uncertainty dummy (26). In the third and fourth columns uncertainty is measured based on the cross-sectional standard deviations, and in the fifth and sixth columns it is measured using the Jurado et al. (2015) uncertainty index. Standard errors, in parentheses, are Newey and West (1994) with automatic lag selection and, for individual-level regressions, clustered by forecaster and time.

***Significant at the 1 percent level. **Significant at the 5 percent level. *Significant at the 10 percent level.

$R_{t+h|t,t-1}$ to vary over time. While this time variation does not affect the pure Bayesian forecast, i.e. when $\theta = 0$ in equation (27), it plays a crucial role under Smooth DE. Specifically, the effective overreaction $\widetilde{\theta}_{t+h|t,t-1}$ increases with this ratio and thus influences forecasts. In particular, $\widetilde{\theta}_{t+h|t,t-1}$ becomes large when current uncertainty is high ($\sigma_{\eta,t} = \sigma^H$) relative to past uncertainty ($\sigma_{\eta,t-1} = \sigma^L$). Intuitively, although the signal $s_{i,t}$ always reduces agents' uncertainty about y_{t+h} compared to prior beliefs, greater residual uncertainty from the temporary component η_{t+h} provides more room for representativeness to amplify overreaction.

We summarize the main results and provide details in Appendix G. First, under Smooth DE, the cross-sectional standard deviation of forecasts $\mathbb{E}^{\theta}_{i,t}[y_{t+h}]$ is larger when $\sigma_{\eta,t} = \sigma^H$ than when $\sigma_{\eta,t} = \sigma^L$, because forecasters overreact more to their private signals $s_{i,t}$ in high-uncertainty environments. This result supports our empirical strategy of identifying high-uncertainty periods using the cross-sectional dispersion of forecasts. In contrast, this dispersion remains constant under both RE and standard DE. Second, the consensus regression coefficient β_C is non-negative as long as forecasters' overreaction to their signals is not "too large." In such cases, agents sufficiently discount their noisy signals, resulting in consensus forecasts that exhibit rigidity. Third, while individual-level forecast error predictability is negative ($\beta_I < 0$) under both standard DE and Smooth DE, it becomes more negative during periods of high uncertainty ($\delta_I < 0$) only under Smooth DE, while under standard DE, predictability remains unchanged ($\delta_I = 0$).

The predictability coefficient from the individual-level regression (25) is given by

$$\beta_I = \frac{\mathbb{COV}\left(y_{t+h} - \mathbb{E}^{\theta}_{i,t}[y_{t+h}], \mathbb{E}^{\theta}_{i,t}[y_{t+h}] - \mathbb{E}^{\theta}_{i,t-1}[y_{t+h}]\right)}{\mathbb{V}\left(\mathbb{E}^{\theta}_{i,t}[y_{t+h}] - \mathbb{E}^{\theta}_{i,t-1}[y_{t+h}]\right)}. \tag{29}$$

Under Smooth DE, the variance and the covariance are time-varying, given by

$$\frac{\mathbb{COV}_t\left(y_{t+h} - \mathbb{E}^{\theta}_{i,t}[y_{t+h}], \mathbb{E}^{\theta}_{i,t}[y_{t+h}] - \mathbb{E}^{\theta}_{i,t-1}[y_{t+h}]\right)}{\mathbb{V}_t\left(\mathbb{E}^{\theta}_{i,t}[y_{t+h}] - \mathbb{E}^{\theta}_{i,t-1}[y_{t+h}]\right)} = -\frac{\widetilde{\theta}_{t+h|t,t-1}\left(1 + \widetilde{\theta}_{t+h|t,t-1}\right)}{\left(1 + \widetilde{\theta}_{t+h|t,t-1}\right)^2 + \left(\rho\widetilde{\theta}_{t+h|t-1,t-2}\right)^2}, \tag{30}$$

which is negative for $\widetilde{\theta}_{t+h|t,t-1} > 0$. Since the unconditional predictability (29) is a weighted average of the conditional predictability (30), we have $\beta_I < 0$. Next, (30) decreases as $\widetilde{\theta}_{t+h|t,t-1}$ increases. Thus the conditional predictability is more negative during periods of high uncertainty ($\delta_I < 0$). Under standard DE, in contrast, the degree of overreaction and the conditional predictability does not vary across states ($\delta_I = 0$).

5 Business cycle implications

We illustrate the business cycle implications of Smooth DE in a parsimonious RBC model with time-varying uncertainty. We argue that Smooth DE emerges as a novel behavioral propagation and amplification mechanism for time-varying uncertainty. We first show that the model replicates, without relying on additional frictions, several salient features of the data thanks to the state-dependent overreaction: (1) *asymmetry* (recessions are sharper than expansions), (2) *countercyclical macro volatility* (time-series variances of macroeconomic variables rise in recessions), and (3) *countercyclical micro volatility* (cross-sectional variances of microeconomic variables rise in recessions). In Appendix H, we provide a general solution algorithm to solve a broad class of models with heteroskedasticity under Smooth DE.

5.1 The model

To isolate the role of Smooth DE as a propagation mechanism, we abstract from conventional frictions in the uncertainty shock literature, such as adjustment costs (Bloom (2009), Bloom et al. (2018)) and sticky prices (Fernández-Villaverde et al. (2015), Basu and Bundick (2017), Bianchi et al. (2023)). The economy consists of a continuum of islands $i \in [0, 1]$. In each island i, an agent has the per-period utility function

$$U(c_{i,t}, h_{i,t}) = \frac{c_{i,t}^{1-\gamma}}{1-\gamma} - \beta \frac{h_{i,t}^{1+\eta}}{1+\eta},$$

where $c_{i,t}$ is consumption, $h_{i,t}$ is the amount of hours worked, γ is the coefficient of relative risk aversion, and η is the inverse of the Frisch labor elasticity. We simplify the algebra below by rescaling the disutility of labor with the discount factor β.

Output in each island is produced as $y_{i,t} = z_{i,t} h_{i,t-1}$. The labor input is chosen before the random realization of productivity $z_{i,t}$ is known. The island resource constraint is $c_{i,t} = y_{i,t}$. We obtain aggregate variables by integrating over the corresponding island variables:

$$H_t = \int_0^1 h_{i,t} di, \quad Y_t = \int_0^1 y_{i,t} di, \quad C_t = \int_0^1 c_{i,t} di$$

The island productivity $z_{i,t+1}$ is the sum of aggregate and idiosyncratic components:

$$\ln z_{i,t+1} = A_{t+1} + a_{i,t+1}.$$

The economy-wide TFP shock A_{t+1} is common across all islands and follows the process

$$A_{t+1} = \rho_A A_t + u_{A,t+1}, \quad u_{A,t+1} \sim i.i.d. N(0, \sigma_A^2).$$

The idiosyncratic TFP $a_{i,t+1}$ is instead specific to island i, and it is composed of a predictable i.i.d. component $\xi_{i,t} \sim N(0, \sigma_\xi^2)$, known one-period-in-advance, and an unpredictable component $u_{a,i,t+1} \sim N(0, \sigma_{a,t}^2)$ realized at $t+1$: $a_{i,t+1} = \xi_{i,t} + u_{a,i,t+1}$. Following Bloom et al. (2018), we assume that the volatility $\sigma_{a,t}$ is time-varying and negatively correlated with the economy-wide TFP. As we describe in Section 5.4, $\sigma_{a,t}$ increases when there is a negative innovation to the economy-wide TFP, and vice versa. We use $\sigma_{a,t}$ to denote the volatility of the period $t+1$ innovation to reflect the assumption that the volatility of the next period's innovation is known one-period-in-advance. In times of low aggregate TFP and high uncertainty, the predictable component $\xi_{i,t}$ reduces uncertainty about $a_{i,t+1}$ less than in times of low uncertainty, as it accounts for a smaller share of the total idiosyncratic variance.

The known time-varying volatility $\sigma_{a,t}$ of the unpredictable idiosyncratic component $u_{a,i,t+1}$ is central to the Smooth DE mechanism. It plays a role analogous to the time-varying volatility of the temporary component in our survey filtering problem (Section 4). In both settings, this feature allows us to isolate the Smooth DE channel: just as the constant precision of private signals in the filtering model attributed changes in overreaction to time-varying uncertainty, here the perfect knowledge agents have about current A_t and their specific $\xi_{i,t}$ ensures that variations in overreaction are driven by time-varying uncertainty. To further isolate the impact of the time-varying volatility $\sigma_{a,t}$ on overreaction and its consequences, we assume that the volatility of the predictable component, $\xi_{i,t}$, is constant. This implies that the cross-sectional dispersion of labor is driven only by the news effect of the uncertainty shock. If we were to relax the assumption of constant volatility of $\xi_{i,t}$, the cross-sectional dispersion would also depend on its realized volatility, but none of the main qualitative properties of the model would change.

5.2 Rational Expectations solution

We first characterize the equilibrium under RE. The island i agent's Bellman equation is

$$\mathcal{V}(h_{i,t-1}, z_{i,t}) = \max_{h_{i,t}} \left\{ U(c_{i,t}, h_{i,t}) + \beta \mathbb{E}_{i,t} \left[\mathcal{V}(h_{i,t}, z_{i,t+1}) \right] \right\}.$$

Combining the first order condition for labor with the envelope condition, we obtain

$$(h_{i,t})^\eta = \mathbb{E}_{i,t} \left[(c_{i,t+1})^{-\gamma} z_{i,t+1} \right]. \tag{31}$$

The optimality condition equates the current marginal disutility of working with its expected benefit. The latter is given by the marginal product of labor weighted by the marginal utility of consumption. We log-linearize the condition and use the method of undetermined coefficients to obtain the RE solution. Hats denote log-deviations from the steady state.

Proposition 2 *The equilibrium under RE is given as follows:*

1. *Individual hours worked are given by* $\widehat{h}_{i,t} = \varepsilon \left[\rho_A A_t + \xi_{i,t}\right]$*, where* $\varepsilon = \frac{1-\gamma}{\eta+\gamma}$*. Equilibrium output and consumption follow immediately as* $\widehat{y}_{i,t} = A_t + a_{i,t} + \widehat{h}_{i,t-1} = \widehat{c}_{i,t}$*.*

2. *Equilibrium aggregate variables are given by* $\widehat{H}_t = \varepsilon \rho_A A_t, \quad \widehat{Y}_t = A_t + \widehat{H}_{t-1} = \widehat{C}_t$

Proof. See Appendix I. ■

The response of individual and aggregate hours to news about expected economy-wide productivity $\rho_A A_t$ and island-specific productivity $\xi_{i,t}$ is affected by the intertemporal elasticity of consumption (IES). When the IES is large enough ($\gamma^{-1} > 1$), an increase in expected productivity raises hours ($\varepsilon > 0$). In that case, the intertemporal substitution effect dominates the wealth effect that would lower hours through the effect on marginal utility.

The next proposition characterizes the cross-sectional variance under RE.

Proposition 3 *The cross-sectional variance of hours worked is given by*

$$\int_0^1 \left(\widehat{h}_{i,t} - \widehat{H}_t\right)^2 di = \varepsilon^2 \sigma_\xi^2,$$

and is constant over the business cycle. The cross-sectional variances of output $y_{i,t}$ *and consumption* $c_{i,t}$ *are increasing in the volatility* $\sigma^2_{a,t-1}$ *of the idiosyncratic TFP.*

Proof. See Appendix I. ■

Under RE, the cross-sectional variance of hours stays constant over the business cycle. This is because once the model is linearized, the news effect of changes in uncertainty is muted under RE. The cross-sectional variances of output and consumption are instead mechanically affected by $\sigma^2_{a,t-1}$ because of the change in realized volatility.

5.3 Smooth DE solution

We now solve the model under Smooth DE. We consider the case of distant memory in which agents' memory recall is based on a more distant past, rather than just the immediate past. This means that the reference group is based on the information set available $J > 1$ periods ago. Bianchi et al. (2024) find that in standard models, distant memory can account for salient features of data, such as persistence and repeated boom-bust cycles. Under distant memory, a time-inconsistency problem arises due to the failure of the law of iterated

expectations. Bianchi et al. (2024) address this issue by adopting the *naïveté* approach (O'Donoghue and Rabin (1999)), which we follow here. Under this approach, the agent fails to take into account that her preferences are time-inconsistent and thinks that in the future she will make choices under perfect memory recall, or RE. However, when the future arrives, the agent ends up changing behavior and being again subject to her imperfect memory recall.

Let θ-superscripts and RE-superscripts denote equilibrium Smooth DE choices and choices under a RE policy function, respectively. The island i agent's Bellman equation is

$$\max_{h_{i,t}^{\theta}} \left\{ U(c_{i,t}^{\theta}, h_{i,t}^{\theta}) + \beta \mathbb{E}_{i,t}^{\theta} \left[\mathcal{V}(h_{i,t}^{\theta}, z_{i,t+1}) \right] \right\},$$

subject to $y_{i,t}^{\theta} = z_{i,t} h_{i,t-1}^{\theta}$ and $c_{i,t}^{\theta} = y_{i,t}^{\theta}$. The continuation value is given by

$$\mathcal{V}(h_{i,t-1}^{\theta}, z_{i,t}) = \max_{h_{i,t}^{RE}} \left\{ U(c_{i,t}^{RE}, h_{i,t}^{RE}) + \beta \mathbb{E}_{i,t} \left[\mathcal{V}(h_{i,t}^{RE}, z_{i,t+1}) \right] \right\},$$

subject to $y_{i,t}^{RE} = z_{i,t} h_{i,t-1}^{\theta}$ and $c_{i,t}^{RE} = y_{i,t}^{RE}$.

Similar to the RE problem, the agent optimally equates the marginal disutility of labor with its expected benefit, except that the benefit is evaluated under Smooth DE:

$$\left(h_{i,t}^{\theta}\right)^{\eta} = \mathbb{E}_{i,t}^{\theta} \left[\left(c_{i,t+1}^{RE}\right)^{-\gamma} z_{i,t+1} \right]. \tag{32}$$

Proposition 4 *The equilibrium under Smooth DE is given as follows:*

1. *Individual hours worked are given by*

$$\begin{aligned} \widehat{h}_{i,t}^{\theta} &= \varepsilon \left[\rho_A A_t + \xi_{i,t} \right] \\ &\quad + \frac{\widetilde{\theta}_{t+h|t,t-J} \eta}{\eta + \left(1 + \widetilde{\theta}_{t+h|t,t-J}\right) \gamma} \varepsilon \left[\rho_A \mathbb{N}_{t-J,t} \left[A_t \right] + \xi_{i,t} \right], \end{aligned} \tag{33}$$

   where ε is given in Proposition 2 and $\mathbb{N}_{t-J,t}\left[A_t\right] \equiv A_t - \mathbb{E}_{t-J}\left[A_t\right]$ represents the news in A_t, compared to past expectation. Equilibrium output and consumption follow as

$$\widehat{y}_{i,t}^{\theta} = A_t + a_{i,t} + \widehat{h}_{i,t-1}^{\theta} = \widehat{c}_{i,t}^{\theta}. \tag{34}$$

2. *The effective overreaction $\widetilde{\theta}_{t+1|t,t-J}$ is given by*

$$\widetilde{\theta}_{t+1|t,t-J} = \frac{R_{t+1|t,t-J}\theta}{1 + \left(1 - R_{t+1|t,t-J}\right)\theta}, \tag{35}$$

where

$$R_{t+1|t,t-J} = \frac{\mathbb{V}_{i,t}\left(-\gamma \widehat{c}^{RE}_{i,t+1} + A_{t+1} + a_{i,t+1}\right)}{\mathbb{V}_{i,t-J}\left(-\gamma \widehat{c}^{RE}_{i,t+1} + A_{t+1} + a_{i,t+1}\right)}. \tag{36}$$

3. Equilibrium aggregate variables are given by

$$\begin{aligned} \widehat{H}^{\theta}_{t} &= \varepsilon \rho_A A_t + \frac{\widetilde{\theta}_{t+1|t,t-J}\eta}{\eta + \left(1 + \widetilde{\theta}_{t+1|t,t-J}\right)\gamma} \varepsilon \rho_A \mathbb{N}_{t-J,t}\left[A_t\right] \\ \widehat{Y}^{\theta}_{t} &= A_t + \widehat{H}^{\theta}_{t-1} = \widehat{C}^{\theta}_{t} \end{aligned} \tag{37}$$

Proof. See Appendix I. ■

Consider the policy function for individual hours $h^{\theta}_{i,t}$. The first line of (33) is identical to the RE policy function. The second line captures the overreaction to news about future productivity. The news, in turn, has two components: a revision in expectations for aggregate productivity and the island specific anticipated component $\xi_{i,t}$. Since $\xi_{i,t}$ is i.i.d., the surprise is $\xi_{i,t}$ itself. Following a positive surprise to the economy-wide TFP A_t, Smooth DE agents become over-optimistic about the future benefit of working, and hence decide to work more in the next period if the substitution effect is strong enough ($\gamma^{-1} > 1$). From (34) and (37), individual and aggregate output and consumption also overreact when individual hours overreact to an aggregate TFP shock.

The coefficient controlling overreaction, $\frac{\widetilde{\theta}_{t+1|t,t-J}\eta}{\eta+\left(1+\widetilde{\theta}_{t+1|t,t-J}\right)\gamma}$, is increasing in $\widetilde{\theta}_{t+1|t,t-J}$. This, in turn, is positively related to $R_{t+1|t,t-J}$, given by (36): the ratio between the current uncertainty about the marginal benefit of labor and the uncertainty perceived at period $t-J$. The effective overreaction $\widetilde{\theta}_{t+1|t,t-J}$ is therefore an equilibrium object. Because agents form expectations about endogenous variables like future consumption, the variance ratio $R_{t+1|t,t-J}$ in equation (36) depends on the stochastic properties of these variables. Consequently, the degree of overreaction depends on the structure of the model and becomes policy dependent.

The expressions (35) and (36) in Proposition 4 suggest that an increase in uncertainty about future idiosyncratic productivity could raise $\widetilde{\theta}_{t+h|t,t-J}$ and, in turn, the overreaction to news.[6] The Proposition below indeed confirms that this is the case.

[6]In the current model, we must take a stance on how agents deal with time-varying volatility when forming expectations. There are two approaches to compute the conditional variance at $t-J$ in (36) that preserve normality of the Smooth DE density. The first approach consists of making an "anticipated utility" assumption (Kreps (1998)). In this case, agents' uncertainty depends on the volatility at the time of the forecast, disregarding the possibility of volatility changes. The second approach consists of assuming that agents take into account the possibility of volatility changes, but that memory retrieves a Normal approximation of the resulting mixture of Normal's. We adopt the first approach, as it is arguably more consistent with the naïveté assumption and the general motivation of DE as a mental heuristic.

Proposition 5 *An increase in the volatility $\sigma_{a,t}^2$ of idiosyncratic TFP raises the effective overreaction $\widetilde{\theta}_{t+1|t,t-J}$.*

Proof. See Appendix I. ■

There are two important implications of this proposition. First, *the business cycle is asymmetric*, even if the underlying shocks are symmetric. A positive TFP shock lowers uncertainty and, as result, overreaction. In contrast, a negative TFP shock raises uncertainty and, as result, overreaction. Thus, a drop in economic activity in response to a negative shock is sharper than an expansion in response to a symmetric positive shock. Second, *macroeconomic volatility is countercyclical*. During expansions uncertainty and overreaction are low while in recessions agents overreact more to economy-wide shocks.

State-dependent overreaction also implies *countercyclical microeconomic volatility*:

Proposition 6 *The cross-sectional variance of hours worked is given by*

$$\int_0^1 \left(\widehat{h}_{i,t}^{\theta} - \widehat{H}_t^{\theta}\right)^2 di = \left[\frac{(1+\widetilde{\theta}_{t+1|t,t-J})(1-\gamma)}{\eta + (1+\widetilde{\theta}_{t+1|t,t-J})\gamma}\right]^2 \sigma_{\xi}^2,$$

and is increasing in $\widetilde{\theta}_{t+1|t,t-J}$ and, thus, in the volatility $\sigma_{a,t}^2$ of the idiosyncratic TFP. The cross-sectional variances of output $y_{i,t}$ and consumption $c_{i,t}$ are similarly increasing in the volatility $\sigma_{a,t-1}^2$ of the idiosyncratic TFP.

Proof. See Appendix I. ■

As uncertainty increases, the overreactions to the predictable component of idiosyncratic TFP and the future benefit of labor, captured in the $\left[\frac{(1+\widetilde{\theta}_{t+1|t,t-J})(1-\gamma)}{\eta+(1+\widetilde{\theta}_{t+1|t,t-J})\gamma}\right]^2$ term, rise. Hence, an increase in uncertainty about idiosyncratic TFP raises the cross-sectional variances of individual choices despite the fact that the predictable component of TFP is homoskedastic.

A visual illustration. Figure 6 provides a visual illustration of the key theoretical results, contrasting the equilibrium functions under Smooth DE with the standard DE and RE. The left panel plots aggregate hours, $\widehat{H}_t$, as a function of the economy-wide TFP innovation $u_{A,t}$. First, both DE and Smooth DE feature steeper slopes than RE, indicating that the DE and Smooth DE policies are more responsive to shocks due to overreaction. Second, the RE and standard DE models feature constant slopes and hence constant responsiveness, while the equilibrium function under Smooth DE is piecewise linear because the sensitivity of aggregate hours to the economy-wide TFP shock is *state-dependent*. As formalized in Proposition 5, the responsiveness to the TFP shock is strongest for large negative innovations (orange dashed line), which under Smooth DE triggers a high-uncertainty state and thus a high effective overreaction. Conversely, the responsiveness is weakest (blue dot-dashed line) for large positive shocks, which are associated with low uncertainty. When uncertainty

Figure 6: Asymmetry and countercyclical volatility

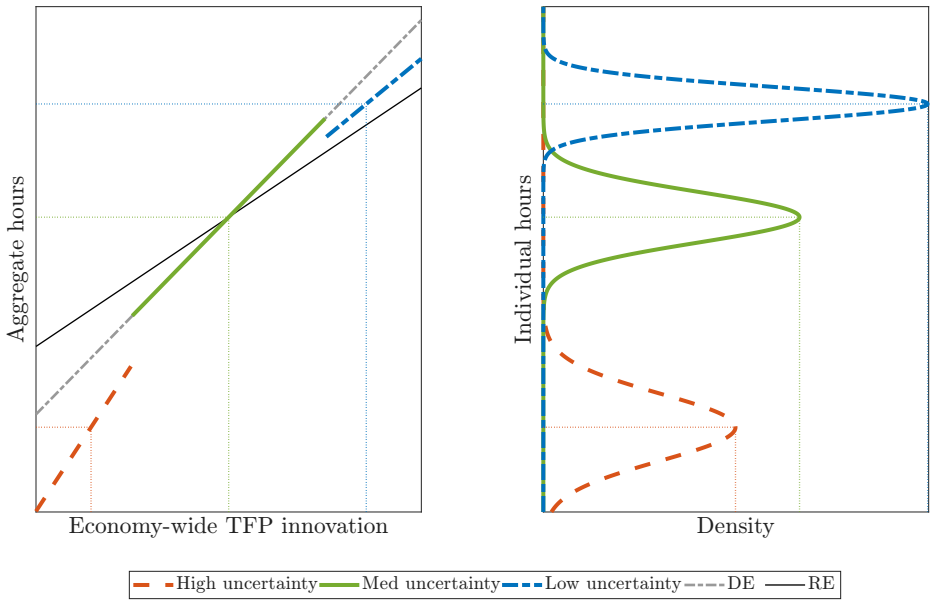

Notes: The left panel plots the policy function for aggregate hours ($\widehat{H}_t$) in response to an economy-wide TFP innovation ($u_{A,t}$) for the Smooth DE (piecewise, in color), standard DE (dashed gray), and RE (solid black) models. The right panel shows the corresponding cross-sectional distribution of individual hours ($\widehat{h}_{i,t}$) under Smooth DE for three illustrative cases. The distributions correspond to a high-uncertainty state triggered by a negative shock (orange dashed line), a medium-uncertainty state (green solid line), and a low-uncertainty state triggered by a positive shock (blue dot-dashed line).

is at the steady-state level, the DE and Smooth DE responses (green solid line) coincide. This state-dependent responsiveness demonstrates the source of the business cycle asymmetry generated by the model: an economy-wide TFP innovation of the same size generates a stronger aggregate hours response for a negative shock than for a positive shock.

The right panel visualizes the implications for the cross-sectional distribution of individual hours under Smooth DE for the three illustrative aggregate shocks. The mean of each distribution corresponds to a point on the aggregate equilibrium function in the left panel. The variance of the cross-sectional distribution is also state-dependent. Following a large negative shock, the increase in uncertainty generates stronger overreaction to the island-specific news, leading to more dispersed choices for hours. The opposite occurs following a large positive shock that lowers uncertainty. This graphically demonstrates the model's key prediction in Proposition 6 of countercyclical micro volatility, despite a homoskedastic process for the predictable component of island-specific TFP shocks.

Policy implication for micro and macro stabilization. Our theory has an im-

portant policy implication. As we saw above, the micro-level volatility and macroeconomic volatility are tightly linked through the state-dependent overreaction controlled by $\widetilde{\theta}_{t+1|t,t-J}$. Thus, a policy that reduces microeconomic uncertainty through, for instance, a redistributive tax policy can also be effective in stabilizing the macroeconomy. To fix ideas, consider a progressive income tax and subsidy scheme where the individual rate $\tau_{i,t}$ is increasing in the realized idiosyncratic productivity level $\tau_{i,t} = \tau a_{i,t}$, where $\tau \geq 0$ is a parameter that controls the progressivity. The island resource constraint is $c_{i,t} + \tau_{i,t} y_{i,t} = y_{i,t}$, so the agent pays a tax ($\tau_{i,t} > 0$) if the realized TFP shock is positive ($a_{i,t} > 0$) and receives a transfer ($\tau_{i,t} < 0$) otherwise. The scheme is budget neutral.

Proposition 7 *A higher progressivity τ is associated with a smaller increase in the effective overreaction $\widetilde{\theta}_{t+1|t,t-J}$ when the volatility $\sigma^2_{a,t}$ of idiosyncratic TFP rises.*

Proof. See Appendix I. ■

This policy is effective because it directly targets the behavioral overreaction channel at the heart of our model. In particular, a progressive tax and transfer system acts as an automatic stabilizer against this behavioral amplification. By making an individual's tax burden dependent on their realized idiosyncratic productivity shock ($a_{i,t}$), the policy provides partial insurance. It dampens the volatility of households' after-tax income and, therefore, their future consumption streams. This insurance directly reduces the conditional variance of the future marginal benefit of labor, the numerator in the conditional variance ratio $R_{t+1|t,t-J}$. By lowering relative uncertainty, the policy mutes the increase in effective overreaction ($\widetilde{\theta}_{t+1|t,t-J}$) that would otherwise occur during a high-uncertainty recession. In essence, the redistributive policy makes agents' degree of overreaction less sensitive to the underlying fluctuations in microeconomic risk. This tames the behavioral response to aggregate shocks, resulting in a milder contraction.

5.4 Quantitative application

We illustrate the quantitative potential of the Smooth DE mechanism in the context of the parsimonious RBC model presented above by examining its dynamics.

Calibration. We calibrate the model to a quarterly frequency. We set the discount factor $\beta = 0.99$, the IES $\gamma^{-1} = 0.25^{-1}$, and $\eta = 0.4$, implying a Frisch elasticity of 2.5.[7] For the economy-wide TFP shock, we set $\rho_A = 0.95$ and $\sigma_A = 0.7/100$. The calibration satisfies the condition $\gamma^{-1} > 1$, so labor increases in response to an increase in expected TFP.

[7]These values of IES and Frisch elasticity allow us to generate realistic labor volatility. Our calibrated model generates the time-series standard deviation of aggregate hours worked of 1.67%. In the data, the standard deviation of total hours worked in the nonfarm business sector (1983:Q1–2019:Q4) is 1.66%.

Table 2: Internally calibrated parameters and targeted moments

Parameters		Targeted moments		
			Data	Model
σ_a	0.022	Realized absolute forecast error	0.143	0.143
σ_ξ	0.027	(Posterior uncertainty)/(ex-ante uncertainty)	0.41	0.41
$\widetilde{\theta}$	1.547	Skewness of aggregate hours	-0.21	-0.21

Notes: The table reports the parameters and their calibrated values as well as the targeted moments.

Consider the time-varying standard deviation $\sigma_{a,t}$ of the idiosyncratic TFP shocks. Using Census micro data, Bloom et al. (2018) and Ilut et al. (2018) find that, during recessions, the dispersion of TFP shocks increases by 13% and 7%, respectively. Motivated by these findings, we assume that a *negative* innovation to economy-wide TFP larger or equal than one standard deviation is associated with a 10% increase in the standard deviation $\sigma_{a,t}$ of the idiosyncratic TFP shocks relative to the steady-state standard deviation σ_a. Conversely, a *positive* economy-wide TFP innovation of the same magnitude is associated with a 10% decrease in the standard deviation $\sigma_{a,t}$ of the idiosyncratic TFP shocks.

We assume that the agent's comparison group is the expectation formed $J = 5$ periods ago. The parameter J mainly determines the persistence of overreaction and the value is consistent with Bianchi et al. (2024), who find that in an estimated structural model the memory weights center around five- and six-quarters-ago expectations.

There are three remaining parameters: the steady-state standard deviations σ_a and σ_ξ of the unpredictable and respectively predictable component of the idiosyncratic TFP shock, and the diagnosticity parameter θ. We calibrate these parameters to match the three empirical moments summarized in Table 2. While multiple model parameters jointly affect these moments, we select the moments that are mostly informative about the corresponding parameters. The first moment, the mean of realized absolute forecast errors, is from Barrero (2022) who uses Survey of Business Uncertainty (SBU), survey data on US managers. Forecast errors are computed by subtracting realized sales growth from t to $t+4$ from managers' forecasts. The model counterpart is obtained by calculating the mean absolute forecast error on the simulated distribution of the realized forecast error $\mathbb{E}^\theta_{i,t}\left[\widehat{y}^{RE}_{i,t+4}\right] - \widehat{y}^\theta_{i,t+4}$.[8] This moment is informative about the steady-state standard deviation of the unpredictable component of idiosyncratic TFP.

[8]Specifically, we generate 100 replications of $T = 200$ time series with $n = 500$ islands. The number of islands roughly matches the number of firms surveyed in the SBU data in Barrero (2022).

Table 3: Untargeted survey moments: overreaction and overconfidence

	(1) $F_t(\Delta y_{i,t+4\|t}) - \Delta y_{i,t+4\|t}$ on $\Delta y_{i,t\|t-1}$	(2) Absolute forecast error: Realized	(3) Absolute forecast error: Subjective	(4) (Subjective)/(Realized)
Data	0.173	0.143	0.023	0.16
	(0.059)	(0.012)	(0.002)	
Model	0.095	0.143	0.017	0.12

Notes: The table reports the coefficient on overextrapolation and realized and subjective mean absolute forecast errors. The data moments are computed by Barrero (2022) using survey data on US managers, with observations employment-weighted and standard errors in parentheses. The model moments are calculated using the simulated data from the Smooth DE model.

The second moment captures the ratio of posterior uncertainty (after seeing $\xi_{i,t}$) relative to the ex-ante uncertainty. David et al. (2016) estimate the ratio to be 41%. The model counterpart is given by $\sigma_a^2/(\sigma_\xi^2 + \sigma_a^2)$. This moment is useful to pin down the standard deviation σ_ξ of the predictable component of idiosyncratic TFP.

Finally, the third moment is the skewness of total hours worked in the nonfarm business sector (1983:Q1–2019:Q4).[9] The negative skewness (-0.21) reflects macroeconomic asymmetry: drops in hours worked are steeper than increases. In our model, a negative economy-wide TFP innovation increases uncertainty $\sigma_{a,t}^2$ and, in turn, the effective overreaction $\widetilde{\theta}_{t+1|t,t-J}$. A positive TFP innovation, in contrast, reduces overreaction. Under Smooth DE, the diagnosticity parameter θ governs the strength of this mechanism to generate asymmetry. Under the RE model and the standard DE model where the overreaction is constant, there is no asymmetry, and the skewness is zero.

The model moments match the empirical moments perfectly. The calibrated σ_a and σ_ξ imply the predictable and unpredictable components' variances are about the same in steady state. The long-run average of effective overreaction $\widetilde{\theta}$, implied by the calibrated value of θ, is 1.54. This value is somewhat larger than Bordalo et al. (2018), Bordalo et al. (2026), and d'Arienzo (2020), who tend to estimate the standard DE diagnosticity parameter around 1, but smaller than the estimate of 1.97 in Bianchi et al. (2024).[10]

Implications for untargeted survey moments. We examine the extent to which our

[9]The empirical skewness of hours increases significantly to -1.75 when we extend the sample until 2022:Q1 to include the 2020 Covid-19 recession. We use the simulated data to compute the skewness of aggregate hours worked in the model. Both simulated and actual time series are HP-filtered with $\lambda = 1600$.

[10]Like Bianchi et al. (2024), our current model features distant memory ($J > 1$). Bianchi et al. (2024) notes that existing estimates are based primarily on models where imperfect memory is assumed to be driven only by the immediate past ($J = 1$), and this assumption changes inference about the diagnosticity parameter.

model can explain untargeted survey evidence on overreaction and overconfidence. We use Barrero (2022)'s survey moments as an external validation, as the study documents both overreaction and overconfidence based on a single dataset (SBU). The first column of Table 3 reports the coefficient from a panel regression where managers' time t forecast of $t+4$ sales growth minus the realization is regressed on the sales growth between quarter $t-1$ and t. The coefficient is positive, indicating that managers' forecasts tend to be excessively optimistic during high-growth periods. The second column contains the realized mean absolute forecast error (MAFE), reported in Table 2, and shown here to facilitate comparison. The third column reports the subjective MAFE, where the hypothetical realizations are drawn from the managers' subjective probability distributions. The fourth column shows that the subjective MAFE is only 16% the size of the realized MAFE, indicating overconfidence.

The model moments are computed by simulation and reported in the third row. First, consider the overextrapolation regression coefficient. The model counterpart is the coefficient on pooled OLS where we regress the Smooth DE four-quarters-ahead forecast error, $\mathbb{E}^{\theta}_{i,t}\left[\widehat{y}^{RE}_{i,t+4}\right] - \widehat{y}^{\theta}_{i,t+4}$, on output growth, $\left[\widehat{y}^{\theta}_{i,t} - \widehat{y}^{\theta}_{i,t-1}\right]$, which proxies for news. The coefficient is positive, but smaller than in the data. The calibrated model understates this coefficient relative to the data because in the model economy-wide shocks are persistent, while island-specific shocks are i.i.d. In contrast to persistent shocks, when shocks are i.i.d., the Smooth DE forecasts are orthogonal to news, so the idiosyncratic shocks push the coefficient toward zero.[11] While we specified idiosyncratic TFP shocks to be i.i.d. for tractability, allowing for persistence would increase the overextrapolation coefficient. Thus, our model provides a conservative lower bound on the macroeconomic effects of Smooth DE.

Next, consider the mean absolute forecast error (MAFE). We calculate this by exploiting the normality of the subjective density. First, we find the agent's subjective variance of output growth using the Smooth DE formula:

$$\mathbb{V}^{\theta}_{i,t}\left[\widehat{y}^{RE}_{i,t+4}\right] = \frac{\mathbb{V}_{i,t}\left[\widehat{y}^{RE}_{i,t+4}\right]}{1 + \left(1 - R_{t+4|t,t-J}\right)\theta}, \tag{38}$$

To convert this subjective variance into the subjective MAFE, we use the standard formula for a Normal distribution's mean absolute deviation, which is its standard deviation multiplied by the constant $\sqrt{2/\pi}$. The resulting subjective MAFE is $\sqrt{2/\pi}\left(\mathbb{V}^{\theta}_{i,t}\left[\widehat{y}^{RE}_{i,t+4}\right]\right)^{\frac{1}{2}}$. The model closely matches the subjective MAFE in the data. The size of the subjective MAFE is 12% of the size of the realized MAFE (fourth column), in line with the survey data's 16%. According to (38), the Smooth DE variance $\mathbb{V}^{\theta}_{i,t}\left[\widehat{y}^{RE}_{i,t+4}\right]$ would be lower than the econome-

[11] This need not be the case when agents forecast endogenous variables in models with slow-moving endogenous states, such as capital. See Bianchi et al. (2024) for details.

Table 4: Countercyclical macro and micro volatility

	(1) Data	(2) Smooth DE	(3) DE	(4) RE
A. Volatility of aggregate labor growth				
(Recessions)/(Expansions)	1.23	1.22	1	1
B. Cross-sectional standard deviation of labor growth				
(Recessions)/(Expansions)	1.16	1.12	1	1

Notes: Panel A reports the ratio of the rolling standard deviation of aggregate labor growth during recessions to the rolling standard deviation during expansions. Panel B reports the ratio of the cross-sectional standard deviation of labor growth during recessions and during expansions.

trician's variance $\mathbb{V}_{i,t}\left[\widehat{y}^{\theta}_{i,t+4}\right]$ due to two factors. First, under *naïveté*, Smooth DE agents perceive future output to follow the RE law of motion $\widehat{y}^{RE}_{i,t+4}$ instead of the equilibrium law of motion $\widehat{y}^{\theta}_{i,t+4}$. Second, under Smooth DE, a reduction of uncertainty associated with a shorter forecast horizon triggers overconfidence, lowering perceived uncertainty. To disentangle these two factors, we calculate the subjective MAFE without the Smooth DE effect. We obtain $\sqrt{2/\pi}\left(\mathbb{V}_{i,t}\left[\widehat{y}^{RE}_{i,t+4}\right]\right)^{\frac{1}{2}} = 0.053$, which is 37% of the size of the realized MAFE. This ratio is more than double the values recovered by the data (16%) and implied by the baseline model (12%). We conclude that the Smooth DE effect is important to account for overconfidence as observed in survey data. Having shown that the model is consistent with key survey evidence, we now turn to its core dynamic mechanism.

Countercyclical macro and micro volatility. In Table 4 we report the model's ability to generate countercyclical macro and micro volatility. To measure time-varying volatility of aggregate hours worked in the data, we compute the rolling window standard deviation as (Ilut et al. (2018)):

$$\sigma_{H,t} = \sqrt{\frac{1}{n_w - 1}\sum_{k=-(n_w-1)/2}^{(n_w-1)/2}\left(\Delta \ln H_{t+k} - \overline{\Delta \ln H_t}\right)^2}, \tag{39}$$

where $\Delta \ln H_{t+k}$ is the log change of total hours worked in the nonfarm business sector from a quarter $t+k-1$ to $t+k$ and $\overline{\Delta \ln H_t} \equiv (1/n_w)\sum_{k=-(n_w-1)/2}^{(n_w-1)/2}\Delta \ln H_{t+k}$. We set the window size $n_w = 3$ and consider the sample 1983:Q1-2019:Q4. The first column reports the measured $\sigma_{H,t}$ during NBER recessions relative to $\sigma_{H,t}$ during NBER expansions. The measured volatility of aggregate labor growth is 23% higher in recessions than in expansions. We then compute the same rolling standard deviation on data simulated from the model. The second column shows that in the model, the aggregate labor growth volatility is 22% higher in

recessions than in expansions. Thus, the Smooth DE model generates countercyclical macro volatility quantitatively in line with the data, even if the volatility of economy-wide shocks is constant. DE and RE models, in contrast, do not generate countercyclical volatility.

Next, consider the micro volatility. In Panel B of Table 4, we report the ratio of the cross-sectional standard deviation of labor growth during recessions to the cross-sectional standard deviation during expansions. The first column shows this ratio from the data, as reported by Ilut et al. (2018), where recessions and expansions are identified based on the NBER labeling. The cross-sectional dispersion is countercyclical: in recessions, it is 16% higher than during expansions. The second column reports the ratio in our model, where we define recessions and expansions as periods when there are one-standard-deviation negative and positive innovations to the economy-wide TFP, respectively. Under Smooth DE the cross-sectional standard deviation of labor growth is 12% higher during recessions than in expansions, so the model explains 75% of the empirical countercyclicality of micro volatility. As shown in the third and the fourth columns, the dispersion is constant under the standard DE model, where the overreaction is constant, and the RE model, where we have $\theta = 0$.

6 Conclusion

We developed a tractable and structural bridge from the representativeness heuristic of Kahneman and Tversky (1972) to the time-series domain. We built on the formalization of representativeness by Gennaioli and Shleifer (2010) and of diagnostic expectations (DE) by Bordalo et al. (2018) to allow for what we call "smooth diagnosticity." Under Smooth DE new information is defined as the difference between the current information set and a previous information set. A critical consequence of this basic approach is that current and past uncertainty interact to determine the intensity of the DE overreaction, but also create the preconditions for novel properties such as over- and underconfidence.

After formally characterizing Smooth DE and its key properties, we leveraged its insights along two substantive directions. First, we showed that Smooth DE implies a joint and parsimonious micro-foundation for key properties of survey data: overreaction to news, stronger overreaction for longer forecast horizons, overconfidence in subjective uncertainty, and a new stylized fact documented in this paper, namely that overreaction is stronger when uncertainty is high. Second, we embedded Smooth DE in a parsimonious RBC model with time-varying uncertainty. This model can account for survey data on overreaction and overconfidence as well as three salient properties of the business cycle: (1) asymmetry, (2) countercyclical macro volatility, and (3) countercyclical micro volatility. We also uncovered a novel policy implication: a redistributive policy that reduces idiosyncratic uncertainty is beneficial for macroeconomic stabilization because it dampens the state-dependent overreaction.

Data Availability Statement

The data and code underlying this research are available on Zenodo at http://doi.org/10.5281/zenodo.19476977.

References

Afrouzi, Hassan, Spencer Y Kwon, Augustin Landier, Yueran Ma, and David Thesmar, "Overreaction in expectations: Evidence and theory," *The Quarterly Journal of Economics*, 2023, *138* (3), 1713–1764.

Altig, David, Jose Maria Barrero, Nicholas Bloom, Steven J Davis, Brent Meyer, and Nicholas Parker, "Surveying business uncertainty," *Journal of Econometrics*, 2022, *231* (1), 282–303.

Angeletos, George-Marios, Zhen Huo, and Karthik A. Sastry, "Imperfect macroeconomic expectations: Evidence and theory," in "NBER Macroeconomics Annual 2020," Vol. 35, The University of Chicago Press, 2021, pp. 1–86.

Augenblick, Ned, Eben Lazarus, and Michael Thaler, "Overinference from weak signals and underinference from strong signals," *The Quarterly Journal of Economics*, 2025, *140* (1), 335–401.

Ba, Cuimin, J Aislinn Bohren, and Alex Imas, "Over-and underreaction to information," *Available at SSRN 4274617*, 2022.

Bachmann, Rüdiger, Steffen Elstner, and Eric R. Sims, "Uncertainty and economic activity: Evidence from business survey data," *American Economic Journal: Macroeconomics*, 2013, *5* (2), 217–249.

Baker, Scott R, Nicholas Bloom, and Stephen J Terry, "Using disasters to estimate the impact of uncertainty," *Review of Economic Studies*, 2024, *91* (2), 720–747.

Barrero, Jose Maria, "The micro and macro of managerial beliefs," *Journal of Financial Economics*, 2022, *143* (2), 640–667.

Basu, Susanto and Brent Bundick, "Uncertainty shocks in a model of effective demand," *Econometrica*, 2017, *85*, 937–958.

Berger, David, Ian Dew-Becker, and Stefano Giglio, "Uncertainty shocks as second-moment news shocks," *Review of Economic Studies*, 2020, *87* (1), 40–76.

Bianchi, Francesco, Cosmin Ilut, and Hikaru Saijo, "Diagnostic business cycles," *The Review of Economic Studies*, 2024, *91* (1), 129–162.

_ , Howard Kung, and Mikhail Tirskikh, "The origins and effects of macroeconomic uncertainty," *Quantitative Economics*, 2023, *14* (3), 855–896.

Bloom, Nicholas, "The impact of uncertainty shocks," *Econometrica*, 2009, *77* (3), 623–685.

_ , "Fluctuations in Uncertainty," *Journal of Economic Perspectives*, 2014, *28* (2), 153–176.

_ **, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry**, "Really Uncertain Business Cycles," *Econometrica*, 2018, *86* (3), 1031–1065.

Bonaglia, Edoardo, Daniele d'Arienzo, Nunzio Fallico, Nicola Gennaioli, and Luigi Iovino, "The Volatility of Inflation Expectations and Interest Rates," 2025. Working Paper, Universit'a Bocconi.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer, "Stereotypes," *The Quarterly Journal of Economics*, 2016, *131* (4), 1753–1794.

_ , _ , _ **, Frederik Schwerter, and Andrei Shleifer**, "Memory and representativeness," *Psychological Review*, 2021, *128* (1), 71–85.

_ **, Nicola Gennaioli, and Andrei Shleifer**, "Diagnostic expectations and credit cycles," *The Journal of Finance*, 2018, *73* (1), 199–227.

_ , _ **, and** _ , "Overreaction and Diagnostic Expectations in Macroeconomics," *Journal of Economic Perspectives*, 2022, *36* (3), 223–44.

_ , _ , _ **, and Stephen J Terry**, "Real credit cycles," *American Economic Review*, 2026, *116* (4), 1274–1308.

_ , _ **, Rafael La Porta, and Andrei Shleifer**, "Diagnostic expectations and stock returns," *The Journal of Finance*, 2019, *74* (6), 2839–2874.

_ , _ **, Rafael La Porta, Matthew OBrien, and Andrei Shleifer**, "Long Term Expectations and Aggregate Fluctuations," *NBER Macro Annual*, 2023.

_ , _ **, Yueran Ma, and Andrei Shleifer**, "Overreaction in macroeconomic expectations," *American Economic Review*, 2020, *110* (9), 2748–82.

Born, Benjamin, Zeno Enders, Gernot J Müller, and Knut Niemann, "Firm expectations about production and prices: Facts, determinants, and effects," *Handbook of Economic Expectations*, 2022.

_ , _ **, Manuel Menkhoff, Gernot J Müller, and Knut Niemann**, "Firm Expectations and News: Micro v Macro," 2025. CEPR Discussion Paper 17768.

Coibion, Olivier and Yuriy Gorodnichenko, "Information rigidity and the expectations formation process: A simple framework and new facts," *American Economic Review*, 2015, *105* (8), 2644–2678.

d'Arienzo, Daniele, "Maturity increasing overreaction and bond market puzzles," 2020. Bocconi University, Working Paper.

David, Joel M., Hugo A. Hopenhayn, and Venky Venkateswaran, "Information, Misallocation, and Aggregate Productivity," *Quarterly Journal of Economics*, 2016, *131* (2), 943—1005.

Fernández-Villaverde, Jésus, Pablo Guerrón-Quintana, Juan F. Rubio-Ramírez, and Martin Uribe, "Risk Matters: The Real Effects of Volatility Shocks," *American Economic Review*, 2011, *101* (6), 2530–2561.

_ , _ , Keith Kuester, and Juan F. Rubio-Ramírez, "Fiscal volatility shocks and economic activity," *American Economic Review*, 2015, *105* (11), 3352–3384.

Gennaioli, Nicola and Andrei Shleifer, "What comes to mind," *The Quarterly Journal of Economics*, 2010, *125* (4), 1399–1433.

_ and _ , *A crisis of beliefs: Investor psychology and financial fragility*, Princeton University Press, 2018.

Halperin, Basil and J Zachary Mazlish, "Overreaction and forecast horizon: longer-term expectations overreact more, shorter-term expectations drive fluctuations," 2025. Working Paper, Department of Economics, University of Oxford.

Hamilton, James D., "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 1989, *57* (2), 357–384.

Ilut, Cosmin L and Martin Schneider, "Ambiguous business cycles," *American Economic Review*, 2014, *104* (8), 2368–99.

_ , Matthias Kehrig, and Martin Schneider, "Slow to Hire, Quick to Fire: Employment Dynamics with Asymmetric Responses to News," *Journal of Political Economy*, 2018, *126* (5), 2011–2071.

Jurado, Kyle, Sydney C. Ludvigson, and Serena Ng, "Measuring Uncertainty," *American Economic Review*, 2015, *105* (3), 1177–1216.

Kahneman, Daniel and Amos Tversky, "Subjective probability: A judgment of representativeness," *Cognitive psychology*, 1972, *3* (3), 430–454.

Kreps, David, "Anticipated utility and dynamic choice," in Donald P. Jacobs, Ehud Kalai, and Morton I. Kamien, eds., *Frontiers of Research in Economic Theory*, Cambridge University Press, 1998.

L'Huillier, Jean-Paul, Sanjay R Singh, and Donghoon Yoo, "Incorporating diagnostic expectations into the New Keynesian framework," *Review of Economic Studies*, 2024, *91* (5), 3013–3046.

McKay, Alisdair and Ricardo Reis, "The Brevity and Violence of Contractions and Expansions," *Journal of Monetary Economics*, 2008, *55* (4), 738–751.

Morley, James and Jeremy Piger, "The Asymmetric Business Cycle," *Review of Economics and Statistics*, 2012, *94* (1), 208–221.

Na, Seunghoon and Donghoon Yoo, "Overreaction and macroeconomic fluctuation of the external balance," *Journal of Monetary Economics*, 2025, *151*, 103750.

Neftci, Salih N., "Are Economic Time Series Asymmetric over the Business Cycle?," *Journal of Political Economy*, 1984, *92* (2), 307–328.

Newey, Whitney K. and Kenneth D. West, "Automatic lag selection in covariance matrix estimation," *Review of Economic Studies*, 1994, *61* (4), 631–653.

O'Donoghue, Ted and Matthew Rabin, "Doing it now or later," *American Economic Review*, 1999, *89* (1), 103–124.

Salgado, Sergio, Fatih Guvenen, and Nicholas Bloom, "Skewed business cycles," Technical Report, NBER Working Paper 26565 2019.

Straub, Ludwig and Robert Ulbricht, "Endogenous uncertainty and credit crunches," *Review of Economic Studies*, 2024, *91* (5), 3085–3115.

Taubinsky, Dmitry, Luigi Butera, Matteo Saccarola, and Chen Lian, "Beliefs About the Economy are Excessively Sensitive to Household-Level Shocks: Evidence from Linked Survey and Administrative Data," 2024. NBER Working Paper 32664.

Tversky, Amos and Daniel Kahneman, "Judgment under uncertainty: Heuristics and Biases," in "Utility, probability, and human decision making" 1975, pp. 141–162.