

Expectations and Learning from Prices

Francesca Bastianello

University of Chicago, Booth School of Business

francesca.bastianello@chicagobooth.edu

Paul Fontanier

Yale University, School of Management

paul.fontanier@yale.edu

April 18, 2024

Abstract

We study mislearning from *equilibrium prices*, and contrast this with mislearning from *exogenous fundamentals*. We micro-found mislearning from prices with a psychologically founded theory of “Partial Equilibrium Thinking” (PET), where traders learn fundamental information from prices, but fail to realize others do so too. PET leads to over-reaction, and upward sloping demand curves, thus contributing to more inelastic markets. The degree of individual-level over-reaction, and the extent of inelasticity varies with the composition of traders, and with the informativeness of new information. More generally, unlike mislearning from fundamentals, mislearning from prices i) generates a two-way feedback between prices and beliefs that can provide an arbitrarily large amount of amplification, and ii) can rationalize both over-reaction and more inelastic markets. The two classes of biases are not mutually exclusive. Instead, they interact in very natural ways, and mislearning from prices can vastly amplify mislearning from fundamentals.

Keywords: Behavioral, Expectations, Inelastic Markets, Learning, Subjective Model

JEL Codes: D83, G40, G41

1 Introduction

The idea that financial markets aggregate dispersed information is central to the Efficient Market Hypothesis, and can be traced back to [Hayek \(1945\)](#). In a world where each agent receives private news about the future payoff of a risky asset and trades on it, the equilibrium price should aggregate all relevant information. However, correctly inferring this information from prices is far from straightforward, as this requires agents to perfectly understand what equilibrium forces generate the price they observe. The Rational Expectations Equilibrium (REE) achieves this level of understanding by assuming common knowledge of rationality. Survey and experimental evidence has however strongly rejected this assumption, by instead finding that misinference is prevalent in games with strategic thinking ([Crawford et al. 2013](#)).

In this paper, we study the asset pricing implications of mislearning from prices. To do so, we first develop a psychologically-founded theory of “Partial Equilibrium Thinking” (PET), where traders misinfer information from prices because they fail to realize the general equilibrium consequences of their actions. Second, we consider a broader class of biases, and study how *mislearning from equilibrium prices* differs and interacts with *mislearning from exogenous fundamentals*.

We show that partial equilibrium thinking provides a micro-foundation for over-reaction to news, where uninformed traders have upward sloping demand curves. This contributes to more *inelastic* markets, and allows us to speak to recent evidence that shows that investors do not adjust their holdings much in response to price changes ([Kojien and Yogo 2019](#), [Gabaix and Kojien 2022](#)). Moreover, the degree of individual level over-reaction and the extent of inelastic demands are environment dependent, and both are more pronounced when there are more uninformed traders, and when the informativeness of news is low.

Turning to the distinction between mislearning from prices and mislearning from fundamentals, we show that these two classes of biases have different qualitative and quantitative implications. In terms of qualitative implications, only with learning from prices can over-reaction contribute to more inelastic rather than more elastic markets. In terms of quantitative implications, unlike mislearning from fundamentals, mislearning from prices generates a two-way feedback between prices and beliefs that can lead to an arbitrarily large amount

of amplification, even for a fixed bias and a bounded signal.

To better understand what we mean by partial equilibrium thinking, consider a trader who sees the price of Apple Inc.’s stock rise, but does not know what caused this price change. They may think that some informed traders in the market received positive news about Apple and decided to buy, pushing up its price. Given this thought process, our uninformed trader infers positive news about Apple, and also decides to buy. However, traders who think in this way fail to realize that every other uninformed trader in the market may be thinking and behaving just like them. Therefore, while part of the price rise they observe is due to the buying pressure of informed agents trading on new information, part of it may be due to the buying pressure of other uninformed traders. If instead uninformed traders attribute the whole price rise to new information alone, they infer news which is better than in reality.

This notion of partial equilibrium thinking builds on extensive evidence that shows that people do indeed tend to make mistakes when learning information from equilibrium outcomes, above and beyond the type of mistakes they might make when learning from exogenous signals (Nagel 1995, Stahl and Wilson 1995, Costa-Gomes et al. 2001; see Crawford et al. 2013 for a survey).¹ While previous papers have considered the financial market implications of not fully using the information in prices (Eyster et al. 2019, Mondria et al. 2022),² we study the implications of over-inference, which in the social learning literature has been widely documented and studied with models of correlation neglect, naïve herding, and K -level thinking (Eyster and Rabin 2008, Eyster and Rabin 2010, and Gagnon-Bartsch and Rabin 2017).³ We build on these models, and introduce this type of bias in a general

¹For example, Esponda and Vespa (2014), Martínez-Marquina et al. (2019), Ngangoué and Weizsäcker (2021) and Keniston et al. (2021) find evidence of under-inference, while Kübler and Weizsäcker (2004), Ziegelmeyer et al. (2013), Penczynski (2017) and Eyster et al. (2018) find evidence of over-inference as agents tend to under-estimate the extent to which others are also learning information from equilibrium outcomes. We consider other types of biases in Section 3, and allow for heterogeneous beliefs in Section 4.

²Eyster et al. (2019) assumes that either traders do not infer information from prices (fully cursed) or that traders act as if they correctly infer information from prices and then only partially use it (partially cursed). Mondria et al. (2022) considers traders who correctly infer information from prices, but then add noise to the signal they inferred. In our model, both these biases can be nested in mislearning from fundamentals, as agents behave as if they had inferred the right information, and then don’t update their beliefs enough in the direction of such news. As a result, there is no two-way feedback between prices and beliefs. Moreover, these models provide a micro-foundation for under-reaction to news, while our model micro-founds over-reaction. Not only can our bias still address excess volume (which is the focus of these other papers), but over-reaction allows us to additionally speak to excess volatility and inelastic markets, as well as survey evidence that investors predominantly over-react to news (Greenwood and Shleifer 2014, Bordalo et al. 2022).

³See Banerjee (1992), Bikhchandani et al. (1992), Chandrasekhar et al. (2020) on rational herding,

equilibrium framework, where prices don't only have an informational role, but also have a market feedback effect role.^{4,5}

Section 2 starts by formalizing the intuition behind the Apple example. We model a continuum of agents who solve a portfolio choice problem between a risky and a riskless asset. A fraction of agents receives news about the future fundamental value of the asset and are thus informed. The remaining agents are uninformed but can infer information from prices. We assume that uninformed traders think in partial equilibrium, and they infer information from prices while failing to realize that all other uninformed traders do so too. This leads them to extract a signal which is biased and more extreme than the true one informed traders hold. The combination of over-reaction and endogenously heterogeneous beliefs then allows us to speak to the asset pricing puzzles of excess volatility, excess trading volume, and return predictability.⁶

Additionally, we offer new predictions that relate to the growing evidence on inelastic markets (Gabaix and Koijen 2022).⁷ Specifically, even though at first sight one might associate the empirically documented insensitivity of holdings to prices with under-reaction, we show that, when traders mislearn information from prices, more inelastic markets can be symptomatic of over-reaction in beliefs, which we need in order to speak to the other asset

DeMarzo et al. (2003), Eyster and Weizsacker (2010), Levy and Razin (2015), Gagnon-Bartsch and Rabin (2017), Enke and Zimmermann (2019) on DeGroot learning and correlation neglect, Bohren (2016), Eyster et al. (2018), Frick et al. (2020), Bohren and Hauser (2021) on misinference, and Esponda and Vespa (2014), Martínez-Marquina et al. (2019), Ngangoué and Weizsäcker (2021), Cohen and Li (2022), Esponda and Vespa (2023), Niederle and Vespa (2023) on the failure of contingent reasoning, among others.

⁴Eyster et al. (2014) present a model of social learning with *congestion costs*, which is a form of market feedback. However, they only consider purely rational learning.

⁵A strand of the social learning literature (e.g., Frick, Iijima and Ishii 2020 or Bohren and Hauser 2021) studies cases where agents are endowed with a set of misspecified models, and among these models they restrict their beliefs to the ones that maximize the fit to the data they observe. The concept of Berk-Nash equilibrium (Esponda and Pouzo 2016) captures learning in those settings. By contrast, our agents consider only one misspecified model of the world, and use this model to update their beliefs about the information possessed by others.

⁶See Shiller (1981) on the excess volatility puzzle, Campbell and Shiller (1988), Cochrane (2011), and Greenwood and Shleifer (2014) on return predictability, and French (2008) on excess trading volume.

⁷Our paper relates to the evidence on the macro elasticity of demand, which studies how agents' holdings change in response to flows from bonds to stocks (Da et al. 2018, Gabaix and Koijen 2022, Hartzmark and Solomon 2022). There is also an extensive body of work estimating the micro elasticity of demand, which studies changes in holdings in response to flows from one stock to another (Shleifer 1986, Chang et al. 2015, Koijen and Yogo 2019, Haddad et al. 2021, Pavlova and Sikorskaya 2023), and these estimates are also much lower than theory implied estimates (Petajisto 2009, Davis et al. 2023). Finally, a more recent literature studies these questions at the factor level as well (Ben-David et al. 2022, Peng and Wang 2023, Li 2022).

pricing puzzles of excess volatility, excess volume and expected return predictability.

To gain intuition on how partial equilibrium thinking speaks to these issues, notice that prices play a dual role in our model. First, prices act as a measure of scarcity. Higher prices make the asset more expensive, leading all traders to want to hold less of it, and pushing towards downward sloping demand curves. Second, prices in our model also have an informational role. Higher prices reflect better fundamentals, leading uninformed traders to want to hold more of the asset, and fuelling a positive feedback effect between prices and beliefs, which pushes towards upward sloping demand curves. Following a price rise(fall), the informational role of prices attenuates agents' desire to decrease(increase) their asset demand due to the asset being more expensive(cheaper), therefore contributing to more inelastic markets.

Moreover, partial equilibrium thinking predicts that the extent of over-reaction and demand inelasticity is environment dependent, with markets exhibiting more over-reaction and greater inelasticity when there are fewer informed traders and when fundamental signals are less informative.⁸ These are environments when the influence on prices of uninformed traders' beliefs is greater. Partial equilibrium thinkers then neglect a greater source of price variation, leading to a stronger bias at the individual level. If the extent of misinference is large enough, even the aggregate demand function for the risky asset may become upward sloping, making the equilibrium unstable. As we approach these unstable regions, the wedge between REE and PET outcomes can grow arbitrarily large, even though at the individual level the nature of the psychological bias itself is fixed.

Section 3 relates our theory of partial equilibrium thinking to the behavioral finance literature on over-reaction surveyed by Barberis (2018).⁹ Many of these biases can be understood as operating through one of two channels in the expectation formation process: mislearning

⁸Our model not only implies that markets exhibit more aggregate level over-reaction when this is the case, but also that the *individual* level of over-reaction varies with these parameters.

⁹We also make a connection to the literature on K -Level thinking and on “dampening general equilibrium” in macroeconomics and finance (Greenwood and Hanson 2015, Angeletos and Lian 2017, García-Schmidt and Woodford 2019, Farhi and Werning 2019 and Iovino and Sergeyev 2020). These papers study setups where K -Level thinkers fail to fully take into account how others respond to fundamental shocks. Importantly, there is no learning from equilibrium outcomes, which is central to our theory. The papers that do consider learning from aggregate outcomes instead assume that agents are rational (Angeletos and Werning 2006, Hellwig et al. 2006, Amador and Weill 2010 and Chahrour and Gaballo 2021).

from prices (Glaeser and Nathanson 2017), or mislearning from fundamentals (Scheinkman and Xiong 2003, Bordalo et al. 2022, Nagel and Xu 2022). To capture these two different classes of biases, we break down the expectations formation process in two steps: whenever traders learn information from prices, we can think of them as first learning fundamental signals from prices; and then processing those fundamental signals to compute posterior beliefs.¹⁰ There are also models that blur the distinction between these two steps (e.g. models of mechanical price and return extrapolation, as in De Long et al. 1990, Hong and Stein 1999, Barberis and Shleifer 2003, Barberis et al. 2018, Jin and Sui 2022). Breaking down the expectation formation process as we do allows us to study both classes of biases separately, before turning to their interaction.

These two classes of biases have different qualitative and quantitative implications. First, the two-way feedback between prices and beliefs embedded in the informational role of prices, and the possibility of unstable equilibria are more general feature of mislearning from prices. Second, mislearning from prices and mislearning from fundamentals predict a different relationship between individual level over-/under-reaction and the extent of demand inelasticity. When traders mislearn from prices, over-reaction is associated with more *inelastic* markets, as uninformed traders engage in positive feedback trading and have upward sloping demand curves (Gabaix et al. 2022). Instead, mislearning from fundamentals is unable to generate upward sloping demand functions for uninformed traders, and the comparative static between the extent of the bias, and demand elasticities is either reversed, or absent altogether. Intuitively, mislearning from fundamentals does not influence the informational role of prices, so that these biases can affect demand elasticities only via distortions in traders' risk-bearing capacities. To the extent that over-reaction from learning from fundamentals is often associated with an increased risk-bearing capacity and agents trading more aggressively against any perceived mispricing, this leads to more *elastic* demand curves, making the empirical evidence on inelastic markets even more puzzling.

¹⁰Anecdotal evidence suggests that traders may indeed (mis)learn fundamental information from prices. For example, Robinhood advises traders to think about their strategy as follows: “[y]our decision to buy or sell a stock should be based on your belief in the value of a company. If the earnings call gives people more confidence in a company, stock prices often go up. If people lose confidence in a company’s performance, stock prices typically go down. Checking if a company’s stock price goes up or down after earnings is a great gut check to see how well the market believes a company is doing” (<https://robinhood.com/us/en/support/articles/using-earnings-a-companys-report-card/>).

Finally, in our framework mislearning from prices and mislearning from fundamentals are not mutually exclusive, consistent with recent evidence that return expectations and earning expectations may be interlinked (Giglio et al. 2021, Beutel and Weber 2022). We show that the interaction between these two classes of biases can arise in very natural ways whenever traders are learning information from prices. For example, misinference from mistakenly assuming the world is rational can vastly amplify biases in learning from fundamentals.

Section 4 considers several extensions, and discusses how our results generalize when agents use heterogeneous models of the world to learn information from prices, and when prices are partially revealing. Section 5 concludes.

2 Model

Setup. We consider a setup similar to Grossman and Stiglitz (1980), after information acquisition costs have been incurred. Agents solve a portfolio choice problem between a risky and a risk-free asset. The risk-free asset is in elastic supply and we normalize its price and its gross rate of return to 1. The risky asset is in fixed supply Z , and has a terminal (next-period) payoff of $v \sim N(\mu_0, \tau_0^{-1})$ per share. We let P be the price of the risky asset.

There is a continuum of measure one of agents with constant absolute risk aversion (CARA) utility over terminal wealth. Agent i 's portfolio choice problem then reduces to (Campbell 2018):

$$\max_{X_i} \left\{ X_i (\mathbb{E}[v|\mathcal{I}_i] - P) - \frac{1}{2} A X_i^2 \mathbb{V}[v|\mathcal{I}_i] \right\} \quad (1)$$

with first-order condition:

$$X_i = \frac{\mathbb{E}[v|\mathcal{I}_i] - P}{A \mathbb{V}[v|\mathcal{I}_i]} \quad (2)$$

where A is the coefficient of absolute risk aversion, X_i captures the number of shares that agent i invests in the risky asset, and \mathcal{I}_i is agent i 's information set. Moreover, we use $\mathbb{E}[\cdot]$ and $\mathbb{V}[\cdot]$ to refer to the mean and variance of a random variable, respectively.

Turning to the information structure, we assume that a fraction ϕ of agents are Informed (I) and receive a common noisy signal $s = v + \epsilon$, where $\epsilon \sim N(0, \tau_s^{-1})$ is independent of v . The remaining fraction $1 - \phi$ of agents are Uninformed (U), and do not directly observe the

realization of s .¹¹ Instead, they can learn information \tilde{s} from prices.^{12,13} Throughout the paper we use $\tilde{\cdot}$ to refer to uninformed traders' beliefs about a variable, so that in this case \tilde{s} is uninformed traders' beliefs of the signal s informed traders received. The setup and the moments which define the conditional distribution of $v|s$ are common knowledge.

Partial Equilibrium Thinking. Understanding what signal \tilde{s} uninformed traders infer from prices is what lies at the core of our theory of partial equilibrium thinking, and this in turn depends on uninformed traders' beliefs of what generates the price that they observe. Under rational expectations, uninformed traders infer the right information from prices because they perfectly understand what generates every price change: they understand that part of it is due to informed traders' response to new information, and part of it is coming from the buying/selling pressure of uninformed traders who learn that information from prices. This level of understanding is achieved via the strong assumption of common knowledge of rationality. We build on experimental evidence that shows that people tend to underestimate the extent to which others also learn information from aggregate outcomes (Kübler and Weizsäcker 2004, Ziegelmeyer et al. 2013, Penczynski 2017, Eyster et al. 2018), and develop a theory of partial equilibrium thinking, where traders misunderstand what generates the price change they observe because they neglect the second source of price variation due to other uninformed traders' learning from prices. Specifically, uninformed PET traders think that they are the only ones learning information from prices, and that all other traders are trading on their private information alone. This leads them to use a misspecified mapping, and to infer biased information from prices, such that $\tilde{s} \neq s$.

¹¹We use this asymmetric information structure to illustrate our notion of PET in the simplest possible framework. In Online Appendix G we show that the intuitions of our model go through even when we allow for a symmetric information structure where all agents receive a private signal.

¹²Prices aggregate dispersed information about fundamentals. Since this information is valuable to uninformed traders (their payoff depends on how accurate their forecast of fundamentals is), they have an incentive to learn that information from prices.

¹³Since prices are fully revealing, uninformed traders extract a signal that is scalar-valued. Online Appendix F extends our model to the case where the supply of the risky asset is stochastic, so that prices are only partially revealing. In this case, agents can only infer a noisy signal of s from prices, which translates into a non-degenerate subjective distribution of \tilde{s} .

Bayesian Updating. Importantly, since prices are fully revealing, uninformed traders think they are inferring the exact signal that informed traders received. Since the conditional distribution of the signal is common knowledge, uninformed traders know the precision of the signal is τ_s .¹⁴ Therefore, given a signal $s_i \in \{s, \tilde{s}\}$, traders obtain their posterior beliefs via Bayesian updating, such that $v|s_i$ is normally distributed with mean $\mathbb{E}[v|s_i] = \frac{\tau_s}{\tau_s + \tau_0} s_i + \frac{\tau_0}{\tau_s + \tau_0} \mu_0$ and variance $\mathbb{V}[v|s_i] = (\tau_s + \tau_0)^{-1}$. In Section 3 we relax this assumption, and we allow for non-rational bayesian updating.

Equilibrium. Given this setup, we can solve for equilibrium outcomes in three steps.

Step 1: True Price Function. We need to determine what truly generates the price agents observe. Given any pair of information sets, $\mathcal{I}_I = \{s\}$ and $\mathcal{I}_U = \{\tilde{s}\}$, we compute agents' asset demand functions and impose market clearing. This gives us the market clearing price P as a function of s and of \tilde{s} .

Step 2: Inference from Prices. We need to determine what uninformed traders think is generating the price they observe. We specify the mapping uninformed agents use to extract information from prices: for any given price P , this mapping returns the signal that uninformed agents extract, \tilde{s} .

Step 3: Equilibrium Price and Extracted Signal. The equilibrium is given by the (P, \tilde{s}) -pair at the intersection of the mappings obtained in Steps 1 and 2, such that uninformed agents' beliefs are consistent with the equilibrium price they observe. Specifically, we define the equilibrium as follows.

Definition 1. *An equilibrium in our economy is a combination of beliefs and prices (\tilde{s}, P) such that:*

- (i) *Agents optimize, according to (2).*

¹⁴We assume that the conditional distribution of the true signal is common knowledge to stay as close as possible to the REE benchmark. In equilibrium, the subjective distribution of $\tilde{s}|v$ will differ from the objective distribution of $s|v$. However, uninformed agents need not realize this as the model is attentionally stable, in the sense of [Gagnon-Bartsch et al. \(2020\)](#). Specifically, notice that uninformed traders think they are using the right model of the world to infer information from prices. Conditional on using the right model, prices are a sufficient statistic for the information they need to learn. This means that they need not pay attention to any other variable, and in particular they need not pay attention to the properties of the conditional distribution of $\tilde{s}|v$. Example 7 in Section 3.2 studies the case where U agents are also misspecified about the precision of the signal received by I agents, such that $\tilde{\tau}_s \neq \tau_s$.

(ii) *The market for the risky asset clears: $\phi X_I + (1 - \phi)X_U = Z$.*

(iii) *Agents' beliefs are consistent with the equilibrium price they observe, given their (mis-specified) model of the world.*

In what follows we solve the model under both rational expectations and under partial equilibrium thinking, and draw a close comparison.

2.1 Step 1: True Price Function

The first step in solving the model requires us to specify the true function that generates the price agents observe, conditional on the exogenous signal s that informed traders receive and on the signal \tilde{s} that uninformed agents learn from prices (and which we solve for as an equilibrium object in Section 2.2). This, in turn, requires us to determine agents' posterior beliefs, so we can compute their asset demands and impose market clearing.

Given the information structure, we know that informed traders receive signal s , and update their beliefs via Bayesian updating:

$$\mathbb{E}[v|\mathcal{I}_I] = \frac{\tau_s}{\tau_s + \tau_0}s + \frac{\tau_0}{\tau_s + \tau_0}\mu_0 \quad (3)$$

$$\mathbb{V}[v|\mathcal{I}_I] = (\tau_s + \tau_0)^{-1} \quad (4)$$

Uninformed traders instead infer signal \tilde{s} from prices, leading to the following posterior beliefs:

$$\mathbb{E}[v|\mathcal{I}_U] = \frac{\tau_s}{\tau_s + \tau_0}\tilde{s} + \frac{\tau_0}{\tau_s + \tau_0}\mu_0 \quad (5)$$

$$\mathbb{V}[v|\mathcal{I}_U] = (\tau_s + \tau_0)^{-1} \quad (6)$$

Substituting these posterior beliefs in traders' demand functions in (2) and imposing market clearing translates into the following price function:

$$P^{True}(s, \tilde{s}) = \frac{\phi\tau_s}{\tau_s + \tau_0}s + \frac{(1 - \phi)\tau_s}{\tau_s + \tau_0}\tilde{s} + \frac{\tau_0}{\tau_s + \tau_0}(\mu_0 - AZ\tau_0^{-1}) \quad (7)$$

which simply reflects that there is a fraction ϕ of traders (the informed) who trade on s with precision τ_s , a fraction $1 - \phi$ of traders (the uninformed) who trade on \tilde{s} with precision τ_s , and a fraction 1 of traders (both informed and uninformed) who trade on the prior with precision τ_0 . This mapping is shown by the solid line in the left panel of Figure 1: given a fixed true signal s , this function provides the market clearing price that arises when I agents trade on s , and all U agents trade on \tilde{s} .

2.2 Step 2: Inference Problem and Perceived Price Function

Next, we need to specify what signal \tilde{s} uninformed traders learn from prices. In what follows, we map out how the inference problem under partial equilibrium thinking differs from the rational benchmark.¹⁵

To solve this inference problem, we ought to understand what uninformed traders think is generating the price they observe. This in turn requires us to understand their beliefs of other agents' posterior beliefs. Both REE and PET uninformed traders think that informed traders receive signal \tilde{s} , and update their beliefs via Bayesian updating:¹⁶

$$\mathbb{E}[v|\tilde{\mathcal{I}}_I] = \frac{\tau_s}{\tau_s + \tau_0} \tilde{s} + \frac{\tau_0}{\tau_s + \tau_0} \mu_0 \quad (8)$$

$$\mathbb{V}[v|\tilde{\mathcal{I}}_I] = (\tau_s + \tau_0)^{-1} \quad (9)$$

REE and PET traders instead differ in their beliefs of other uninformed traders' beliefs. REE traders think that uninformed traders have the same beliefs as informed traders, as they recognize that uninformed traders are able to infer the right information from prices:

$$\mathbb{E}[v|\tilde{\mathcal{I}}_U]_{REE} = \frac{\tau_s}{\tau_s + \tau_0} \tilde{s} + \frac{\tau_0}{\tau_s + \tau_0} \mu_0 \quad (10)$$

$$\mathbb{V}[v|\tilde{\mathcal{I}}_U]_{REE} = (\tau_s + \tau_0)^{-1} \quad (11)$$

Instead, PET traders mistakenly think that other uninformed traders are not learning in-

¹⁵Appendix B further highlights that the inference problem of uninformed agents in the REE involves a complicated fixed-point problem, which is instead absent for uninformed PET traders, who instead solve a much simpler inference problem.

¹⁶Recall that \tilde{s} is uninformed traders' beliefs of the signal informed traders received.

formation from prices, and believe that other uninformed traders are trading on their prior beliefs:

$$\mathbb{E}[v|\tilde{\mathcal{I}}_U]_{PET} = \mu_0 \quad (12)$$

$$\mathbb{V}[v|\tilde{\mathcal{I}}_U]_{PET} = (\tau_0)^{-1} \quad (13)$$

Imposing market clearing then leads REE and PET traders to use different mappings to infer information from prices. Rational uninformed traders invert the following price function:¹⁷

$$P_{REE}^{Mis}(\tilde{s}) = \frac{\tau_s}{\tau_s + \tau_0} \tilde{s} + \frac{\tau_0}{\tau_s + \tau_0} (\mu_0 - AZ\tau_0^{-1}) \quad (14)$$

which reflects that a fraction 1 of traders (both informed and uninformed) trade on \tilde{s} with precision τ_s and on the prior with precision τ_0 . Instead, PET traders think that prices are generated by the following price function:

$$P_{PET}^{Mis}(\tilde{s}) = \frac{\phi\tau_s}{\phi\tau_s + \tau_0} \tilde{s} + \frac{\tau_0}{\phi\tau_s + \tau_0} (\mu_0 - AZ\tau_0^{-1}) \quad (15)$$

which reflects U agents' beliefs that only a fraction ϕ of agents (i.e. only the informed) trade on \tilde{s} with precision τ_s , and a fraction 1 of agents (both the informed and the uninformed) trade on the prior with precision τ_0 .¹⁸ Therefore, relative to the rational benchmark, uninformed PET traders fail to take into account that the fraction $1 - \phi$ of traders who are uninformed also trade on the signal they inferred from prices. The PET mapping in (15) is shown by the dashed line in the left panel of Figure 1: given any price P on the vertical axis, this mapping returns on the horizontal axis the signal \tilde{s} that PET agents extract.^{19,20}

¹⁷Appendix B shows more formally that this is the result of a fixed-point argument, which is instead absent in the derivation of the PET mapping. That allows us to highlight how the PET mapping is cognitively much simpler than the REE counterpart. Online Appendices F and G further considers the case with noisy supply such that prices are only partially revealing. In that case, the fixed-point problem embedded in the REE is even more involved, and using PET inference is once again cognitively much simpler.

¹⁸In order to stay as close as possible to the REE benchmark, here we assume that PET traders know the fraction of informed traders, ϕ . We allow for uninformed traders to hold misspecified beliefs about this quantity ($\tilde{\phi} \neq \phi$) in Section 3: $0 < \tilde{\phi} < \phi$ amplifies our results further, while $\tilde{\phi} > \phi$ dampens them.

¹⁹Because signals are normally distributed, and prices in the misspecified model are linear in the signal, any observed price can be rationalized by some \tilde{s} .

²⁰The PET inference process can also be seen as 2-level thinking. Online Appendix D allows for K -level thinking and shows that it converges to the REE fixed-point as K grows to infinity if and only if $\phi > 1/2$.

Comparison of the PET and REE mappings. Comparing the slope of the PET mapping in (15) to the rational counterpart in (14) is informative of the nature of misinference inherent in partial equilibrium thinking:

$$\frac{\partial P_{PET}^{Mis}(\tilde{s})}{\partial \tilde{s}} = \frac{\phi \tau_s}{\phi \tau_s + \tau_0} < \frac{\tau_s}{\tau_s + \tau_0} = \frac{\partial P_{REE}^{Mis}(\tilde{s})}{\partial \tilde{s}} \quad (16)$$

PET agents think that the price is less responsive to new information than it truly is, as they think that uninformed agents trade less aggressively on new information than they really do. By under-estimating the sensitivity of prices to new information, PET agents attribute too much of the price change they observe to new information and extract a signal that is more extreme than in reality, as in the Apple example in the introduction. The extent of misinference is decreasing both in the fraction of informed traders (ϕ), and in the precision of the signal relative to the prior (τ_s/τ_0). Intuitively, when ϕ and τ_s/τ_0 are lower, uninformed traders have a greater influence on prices, leading partial equilibrium thinkers to neglect a greater source of price variation.

Inference and Demand Curves. Not only does PET lead to over-reaction to news, but it does so by inducing uninformed traders to have upward sloping demand curves. To see why this is the case, we can compare the slope of informed and uninformed traders' demand curves as follows:

$$\frac{\partial X_I}{\partial P} = -\frac{1}{A(\tau_s + \tau_0)^{-1}} < 0 \quad (17)$$

$$\frac{\partial X_U}{\partial P} = -\frac{1}{A(\tau_s + \tau_0)^{-1}} \left(1 - \frac{\tau_s}{\tau_s + \tau_0} \frac{\partial \tilde{s}}{\partial P} \right) = -\frac{1}{A(\tau_s + \tau_0)^{-1}} \left(1 - \frac{1 + \frac{1}{\phi} \frac{\tau_0}{\tau_s}}{1 + \frac{\tau_0}{\tau_s}} \right) > 0 \quad (18)$$

While the slope of informed traders' demand curve is unambiguously negative, the slope of uninformed traders' demand curve has two terms that push in opposite directions in determining its sign. These two terms reflect that prices play a dual role in this general equilibrium environment. First, they have their standard role as a measure of scarcity: higher prices make the asset more expensive, inducing all agents to reduce their asset demand. Second, prices also have an informational role: higher prices reflect better fundamentals,

leading uninformed agents to want to hold more of the asset.

The second equality in (18) shows that with partial equilibrium thinking the informational role of prices dominates, and uninformed traders have upward sloping demand curves. Intuitively, following good news, PET traders think that other uninformed traders are under-reacting to news, meaning that they think that the asset is undervalued, and this turns them into momentum traders. This is in contrast to the rational counterpart, where $\frac{\partial \tilde{s}}{\partial P} = \frac{\tau_s + \tau_0}{\tau_s}$ and $\frac{\partial X_U^{REE}}{\partial P} = 0$, so that uninformed rational traders have perfectly inelastic demand curves, as they think that the asset is always correctly priced.

At the aggregate level, partial equilibrium thinking contributes to more inelastic markets than under rational expectations thus providing an additional channel in understanding the causes of inelastic markets (Gabaix and Koijen 2022): following a price increase, PET's stronger informational role of prices leads to a greater dampening of the reduction in holdings due to the asset becoming more expensive.

Proposition 1 (Elasticity of Demand with Partial Equilibrium Thinking). *Partial equilibrium thinking leads uninformed traders to have upward sloping demand curves, therefore making the aggregate demand for the risky asset more inelastic than under rational expectations. Moreover, markets are more inelastic when there are fewer informed traders in the market, and when the precision of the signal relative to the prior is low.*

Proof. All proofs are in Appendix A, unless stated otherwise. □

Empirical Counterparts. There are a number of empirical findings that relate to our predictions. First, with partial equilibrium thinking uninformed traders' demand elasticity depends on how aggressively they think others trade, which is consistent with the type of behavior modeled and documented in Haddad et al. (2021).²¹ Second, the upward sloping nature of uninformed PET agents' demand curves is consistent with recent evidence on US

²¹Specifically, Haddad et al. (2021) model traders' demand elasticities as being linear in the aggregate demand elasticity. Partial equilibrium thinking provides a micro-foundation for this, with the slope of uninformed traders' demand curves being linear in their *perception* of the slope of the aggregate demand curve. To see why this is the case, notice that we can re-write (18) as: $\frac{\partial X_U}{\partial P} = -\frac{1}{A\bar{V}_I} + \frac{1}{\phi} \frac{\partial \tilde{X}_{TOT}}{\partial P}$, where $\frac{\partial \tilde{X}_{TOT}}{\partial P} = \frac{\phi}{A\bar{V}_I} + \frac{1-\phi}{A\bar{V}_U}$ is PET traders' perception of the slope of the aggregate demand curve with no learning from prices.

households. Specifically, [Gabaix et al. \(2022\)](#) study the portfolio rebalancing behavior of US households and find that less wealthy households act as momentum traders during market downturns, with wealthier households taking the opposite side. If we think of wealth levels as proxying how “informed” agents are likely to be, then this trading behavior is consistent with our model’s prediction that informed (wealthier) agents have more elastic demands than uninformed (less wealthy) agents, with the latter group having upward sloping demand curves. Moreover, [Kojen et al. \(2020\)](#) find time-variation in the weights that investors put on stock characteristics in forming their beliefs. One interpretation for these findings is that the way investors process information is environment dependent. Partial equilibrium thinking provides a belief-based explanation for such environment-dependence, and additionally predicts that this variation (as well as the extent of demand inelasticity) should depend on the composition of traders in the market, and on the informativeness of new information.

To test the predictions in Proposition 1 more directly, we need empirical proxies for the fraction of informed agents and for the informativeness of news. Prior work has found a number of possible candidates (see [Veldkamp \(2023\)](#) for a review). For example, [Gompers and Metrick \(2001\)](#), [Boehmer and Kelley \(2009\)](#) and [Yan and Zhang \(2009\)](#) use the share of institutional investors to proxy for informed traders, while [Laarits and Sammon \(2022\)](#) use the fraction of retail traders as a proxy for uninformed trading. Turning to the precision of new information, [Hong et al. \(2000\)](#) proxy this with the number of analysts covering a given stock, while [Bae et al. \(2008\)](#) uses the precision of forecasts reported in survey data. Understanding how demand inelasticity varies with these quantities, both in the time-series and in the cross-section, would shed further light in understanding potential sources of misinference.

2.3 Step 3: Equilibrium Price and Extracted Signal

The equilibrium price and extracted signal are given by the intersection of the true and misspecified mappings we just derived (as in the left panel of Figure 1), as this is the only point where the signal uninformed agents extract from prices is consistent with the equilibrium price they observe. To see why, suppose instead that we were to start from a price away from the intersection of the true and misspecified mappings, as in the left

panel of Figure 1. Following the arrows, given this price, uninformed traders would use the misspecified mapping to infer a signal from prices. But if all uninformed traders were to trade on that extracted signal, the true mapping shows that the market clearing price would be higher, leading uninformed traders to extract an even higher signal, and so on. Instead, at the intersection, if uninformed traders observe P , they extract signal \tilde{s} , and if all uninformed traders trade on \tilde{s} the price that clears the market is precisely P .

Rational Expectations Equilibrium. Solving for the equilibrium price and extracted signal $(P_{REE}, \tilde{s}_{REE})$ -pair such that $P_{REE} \equiv P^{True}(s, \tilde{s}_{REE}) = P_{REE}^{Mis}(\tilde{s}_{REE})$ in (7) and (14), we find:

$$P_{REE} = \frac{\tau_s}{\tau_s + \tau_0} s + \frac{\tau_0}{\tau_s + \tau_0} (\mu_0 - AZ\tau_0^{-1}) \quad (19)$$

$$\tilde{s}_{REE} = s \quad (20)$$

Comparing (14) to (19) we see that when uninformed traders have rational expectations they use the right mapping to infer information from prices, which in turn allows them to infer the right signal, as shown in (20).

Partial Equilibrium Thinking. Solving for the equilibrium price and extracted signal $(P_{PET}, \tilde{s}_{PET})$ -pair such that $P_{PET} \equiv P^{True}(s, \tilde{s}_{PET}) = P_{PET}^{Mis}(\tilde{s}_{PET})$ in (7) and (15), we find:

$$P_{PET} = \left(\frac{\tau_s}{\tau_s + \tau_0 - \left(\frac{1-\phi}{\phi}\right)^2 \tau_0} \right) s + \left(1 - \frac{\tau_s}{\tau_s + \tau_0 - \left(\frac{1-\phi}{\phi}\right)^2 \tau_0} \right) (\mu_0 - AZ\tau_0^{-1}) \quad (21)$$

$$\tilde{s}_{PET} = \left(1 + \frac{\left(\frac{1-\phi}{\phi^2}\right) \tau_0}{\tau_s + \tau_0 - \left(\frac{1-\phi}{\phi}\right)^2 \tau_0} \right) s + \left(\frac{\left(\frac{1-\phi}{\phi^2}\right) \tau_0}{\tau_s + \tau_0 - \left(\frac{1-\phi}{\phi}\right)^2 \tau_0} \right) (\mu_0 - AZ\tau_0^{-1}) \quad (22)$$

where the denominator in (21) and (22) is positive as long as the aggregate demand for the risky asset is downward sloping, so that the equilibrium is stable:

$$\frac{\partial X_{TOT}}{\partial P} = -\frac{1}{A(\tau_s + \tau_0)^{-1}} \left(1 - \frac{(1-\phi)\tau_s}{\tau_s + \tau_0} \times \frac{\phi\tau_s + \tau_0}{\phi\tau_s} \right) < 0 \quad (23)$$

This, in turn, requires that in the aggregate the informational role of prices is weaker than the scarcity role.

Lemma 1 (Stability). *The PET equilibrium is stable when the aggregate excess demand function for the risky asset is downward sloping. In our setup, this requires:*

$$\phi > \left(1 + \sqrt{1 + \frac{\tau_s}{\tau_0}}\right)^{-1} \quad (24)$$

Therefore, the stability region is increasing both in the fraction of informed agents, ϕ , and in the informativeness of the signal relative to the prior, τ_s/τ_0 .

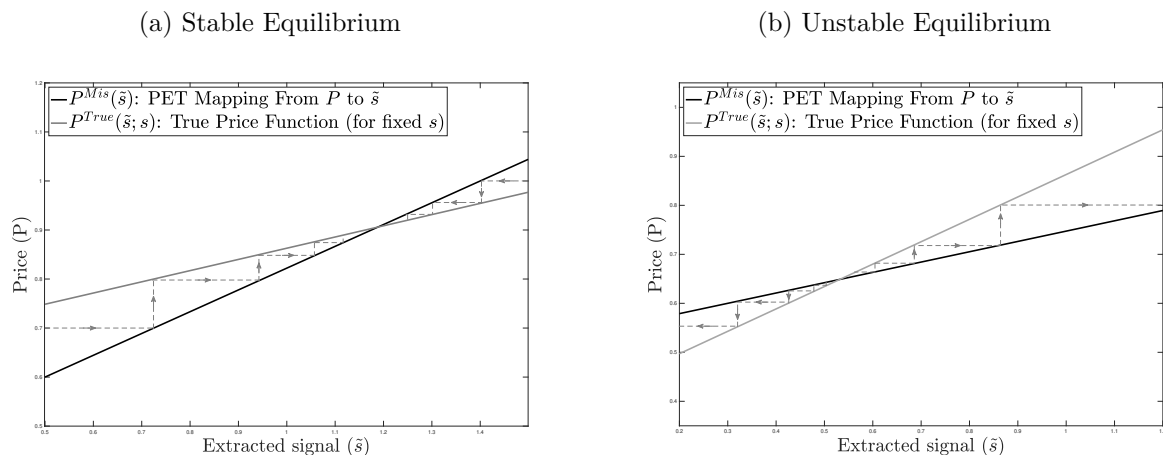
We can illustrate this notion of stability graphically by noticing that the strength of the informational role of prices in (23) is determined by the relative slope of the true and misspecified mappings in (7) and (15) respectively, and that stability simply requires that the true mapping be steeper than the misspecified mapping ($1 - \frac{\partial P^{True}(s, \tilde{s})}{\partial \tilde{s}} / \frac{\partial P^{Mis}(\tilde{s})}{\partial \tilde{s}} > 0$), as in the left panel of Figure 1. Intuitively, the shallower the PET mapping in (15) is, the greater the wedge between the PET and the rational mapping, and the greater the extent of over-reaction, as discussed in Section 2.2. Moreover, the steeper the true mapping in (7) is, the greater is the influence on prices of uninformed traders' biased beliefs. Both these forces contribute to amplifying the informational role of prices. The relative slope of the true and misspecified mappings thus exactly quantifies the strength of the feedback effect between prices and beliefs:

$$\text{strength of feedback effect} = \frac{\partial P^{True}(\tilde{s})}{\partial \tilde{s}} / \frac{\partial P^{Mis}(\tilde{s})}{\partial \tilde{s}} = \frac{(1 - \phi)\tau_s}{\tau_s + \tau_0} \times \frac{\phi\tau_s + \tau_0}{\phi\tau_s} \quad (25)$$

If the true mapping becomes steeper than the misspecified mapping, then the informational role of prices dominates over the scarcity role, leading to an upward sloping aggregate demand curve. The right panel of Figure 1 shows that when this is the case the feedback between prices and beliefs implicit in partial equilibrium thinking becomes so strong that the equilibrium becomes unstable, and if we were to start calling out prices and beliefs away from the steady state, a tâtonnement process would diverge away.

In the rest of the paper, we focus on the properties of stable equilibria, so that the

Figure 1: PET Feedback Effect and Stability. This figure illustrates the feedback effect between prices and beliefs implicit in partial equilibrium thinking, and graphically shows how the equilibrium can be either stable or unstable. The solid black line plots the mapping used by uninformed traders to infer information from prices, while the solid gray line plots the true market clearing price function when informed agents trade on a fixed signal s , and uninformed agents extract a signal \tilde{s} from prices. The intersection of the two mappings gives the PET equilibrium price and extracted signal for a given s . The arrows trace out the *tâtonnement* process that would arise if we started at a point away from the intersection of the two mappings. Panel (a) provides an example of a stable equilibrium, and shows that in this case the *tâtonnement* process converges to the steady state equilibrium. Conversely, Panel (b) depicts a case where the equilibrium is unstable, and shows that in this case the *tâtonnement* process leads prices and beliefs to diverge away from the steady state.



denominator in (21) and (22) is positive, and prices and beliefs are increasing in the terminal payoff. The reason why we focus on stable equilibria is that we are studying the properties of *steady state* outcomes, and empirical evidence suggests that the aggregate excess demand function is indeed downward sloping on average.²² It is however worth highlighting that unstable regions are simply symptomatic of a feedback effect that is so strong that it becomes explosive. In what follows we show that as we approach these unstable regions the strength of the feedback effect between prices and beliefs becomes stronger, and deviations from the rational benchmark can become arbitrarily large.²³ Finally, Section 3 shows that the possible unstable nature of equilibrium outcomes is widespread in misinference models, and

²²As prices deviate farther from fundamentals, informed traders earn higher payoffs, therefore increasing the incentives to be informed. Online Appendix C shows that allowing for endogenous information acquisition (on either the intensive or the extensive margin) plays a stabilizing force, and ensures that the equilibrium is stable as long as entry costs are not too high. Even though endogenous information acquisition contributes to ensuring stability, the equilibrium can still approach unstable regions, so our insights about amplification are robust to this extension.

²³In Bastianello and Fontanier (2023) we develop a dynamic model that exploits both stable and unstable equilibria to explain both normal times market dynamics and the formation of bubbles and crashes.

not just a curiosity of PET.

2.4 Over-reaction, Amplification, and Environment Dependence

Comparing the PET equilibrium price and extracted signal in (21) and (22) to their rational counterparts in (19) and (20), we notice two key properties of PET. First, PET agents extract a biased signal from prices, therefore leading to *endogenously heterogeneous beliefs* without relying on heterogeneous priors, or agree to disagree motives. Second, the type of misinference induced by partial equilibrium thinking provides a micro-foundation for *over-reaction to news*.

Relative to other models of over-reaction, the micro-foundation provided by partial equilibrium thinking has three key features. First, the *size of the bias is environment dependent*, with over-reaction being more accentuated when there are less informed traders in the market, and when the precision of the signal is low relative to the prior. While it is not necessarily surprising that the aggregate level of over-reaction depends on the composition of traders, the fact that the individual level of over-reaction also depends on the composition of traders is a distinguishing feature of our model that can be tested empirically.²⁴ Intuitively, a smaller fraction of informed traders in the market, and a lower precision of the signal relative to the prior both increase the influence on prices of uninformed traders' beliefs. This leads partial equilibrium thinkers to neglect a greater source of price variation, which in turn translates into a stronger bias.

Corollary 1 (Over-reaction and Environment-dependence). *Partial equilibrium thinking provides a micro-foundation for over-reaction to news, where the degree of both individual and aggregate level over-reaction is decreasing in the fraction of informed traders in the market, and in the informativeness of the signal.*

Second, as we approach unstable regions, departures from rationality can be *arbitrarily large*, without relying on extreme signals. This is somewhat striking given that at the individual level the nature of the psychological bias is fixed: regardless of the environment, PET truncates common knowledge of rationality to rationality of second order beliefs.

²⁴For example, in models of over-confidence or diagnostic expectations the extent of individual level over-reaction is assumed to be constant across environments.

Proposition 2 (Amplification). *Assume that $\phi\tau_s/(\tau_0 + \tau_s)$ and $(1 - \phi)\tau_s/(\tau_0 + \tau_s)$ stay bounded above and below by strictly positive quantities. For a fixed true signal, when the PET equilibrium is stable, the wedge between PET and REE equilibrium outcomes (extracted signals, prices, sensitivities to information, expected returns) grows arbitrarily large in the limit where the feedback effect goes to 1 from below:*

$$\frac{(1 - \phi)\tau_s}{\tau_s + \tau_0} \times \frac{\phi\tau_s + \tau_0}{\phi\tau_s} \rightarrow 1^-. \quad (26)$$

Third, as already discussed in Section 2.2, PET generates over-reaction in a way that leads some traders to have upward sloping demand curves, therefore also contributing to more *inelastic markets*.

2.5 Asset Pricing Puzzles

The combination of over-reaction and endogenously heterogeneous beliefs implies that partial equilibrium thinking lends itself to explaining the asset pricing puzzles of excess volatility, excess volume, and expected return predictability, as well as providing an additional source of inelastic demands.

Corollary 2 (PET and Asset Pricing Puzzles). *The stable PET equilibrium lends itself to reconciling the following asset pricing puzzles.*

1. **Over-reaction and Excess Volatility:** *PET delivers over-reaction to news:*

$$\frac{\partial P_{PET}}{\partial s} > \frac{\partial P_{REE}}{\partial s}; \quad (27)$$

2. **Excess Trading Volume:** *For any fixed $s \neq \mu_0 - AZ\tau_0^{-1}$, PET trading volume is strictly positive, and unbounded above:*

$$V_{PET} \rightarrow +\infty \text{ when } \phi \rightarrow \left(\left(1 + \sqrt{1 + \frac{\tau_s}{\tau_0}} \right)^{-1} \right)^+ \quad (28)$$

3. **Expected Return Predictability:** Define expected returns as:

$$\mathbb{E}_i[R] = \mathbb{E}[v|\mathcal{I}_i] - P \quad (29)$$

On average:

- $\frac{\partial \mathbb{E}_U[R]}{\partial s} > 0$: *U agents' expected returns are higher in good times;*
- $\frac{\partial R}{\partial s} < 0$: *realized returns are lower in good times.*

4. **Inelastic Demands:** *PET contributes to more inelastic aggregate demand curves, and induces uninformed traders to have upward sloping demand curves. Moreover, individual elasticities are environment dependent, and vary with the composition of traders in the market, and with the informativeness of signals relative to the prior.*

Over-reaction and excess volatility follow immediately from Proposition 1.²⁵ Heterogeneous beliefs additionally yield excess volume, which can be arbitrarily large for environments that approach unstable regions (when the feedback effect increases towards 1). Turning to expected return predictability, this is also a consequence of over-reaction and heterogeneous beliefs: uninformed traders become over-optimistic in good times, only to be systematically disappointed, as documented empirically in Greenwood and Shleifer (2014).²⁶ Our model has a natural explanation: in bad times investors fail to realize that other investors hold the same beliefs as them, and that the equilibrium price is lower to reflect this. This leads them to expect lower returns following bad news, as is the case for uninformed traders in our model. More generally, these findings are consistent with empirical studies that challenge the conventional wisdom on risk-based explanations for time-varying risk-premia. Partial equilibrium thinking also breaks the classic risk/return trade-off and instead delivers time-varying expected returns that do *not* ultimately compensate investors for bearing risk.²⁷

²⁵De La O and Myers (2021) and Bordalo et al. (2022) use survey data to show that over-reaction in cash-flow expectations can account for Shiller's excess volatility puzzle.

²⁶Notice that the combination of two key ingredients allows us to speak to the empirical patterns of return predictability documented in Greenwood and Shleifer (2014): heterogeneous agents, and over-reaction. One without the other would not be enough in our setup. For example, models of over-reaction with a representative agent are not able to generate any variation in expected returns, unless they allow for the perceived variance of the fundamental to vary: with $\overline{Var}[v]$ constant, homogeneous beliefs imply that $\mathbb{E}[v|\mathcal{I}_i] - \bar{\mathbb{E}}[v] = 0$, so that expected returns in equation (29) must also be constant.

²⁷Moreira and Muir (2017) show that a strategy that manages volatility so as to take *less* risk in recessions

Finally, while the behavioral finance literature has focused on the three asset pricing puzzles of excess volatility, return predictability and excess trading volume, we additionally draw attention to the role of behavioral biases in speaking to the growing evidence on inelastic markets. Partial equilibrium thinking is consistent with both over-reaction and more inelastic markets, and it offers additional predictions on the sources of heterogeneity and environment dependence of demand elasticities.

Moreover, there is a close connection between the extent of excess volatility and inelasticity that partial equilibrium thinking contributes to. To see this, we can define the extent of excess volatility as:

$$\text{ExcVol} \equiv \left(\frac{\partial P_{PET}}{\partial s} \bigg/ \frac{\partial P_{REE}}{\partial s} \right)^2$$

We can then relate the slope of the aggregate demand curve under partial equilibrium thinking and under rational expectations as follows:

$$\frac{\partial X_{TOT}^{PET}}{\partial P} = \frac{1}{\sqrt{\text{ExcVol}}} \frac{\partial X_{TOT}^{REE}}{\partial P}$$

meaning that partial equilibrium thinking contributes to more inelastic markets, by shrinking the slope of the rational aggregate demand curve by a factor equal to the excess standard deviation. This disciplines by how much our theory can jointly explain asset pricing puzzles. For instance, if misinference is responsible for a degree of over-reaction that leads to prices being twice as sensitive to information compared to the rational counterfactual, then misinference also explains a market elasticity halved compared to the rational counterfactual. Given the substantial amount of excess belief volatility compared to rational expectations ([Augenblick and Lazarus 2023](#)), we view our theory as a promising way to jointly explain these puzzles.

We return to this in Section 3.3, where we show that these prediction differentiate PET from conventional models of mislearning from fundamentals.

and crisis still earns high average returns, contrary to what leading models would predict. [Jensen \(2023\)](#) shows that most factors earning a positive realized excess return are actually perceived as safer by analysts.

3 Mislearning from Prices and from Fundamentals

As discussed in the introduction, whenever traders learn information from prices, we can break down the expectation formation process in two steps: first, uninformed traders infer fundamental information \tilde{s} from prices, and second they combine this new information with their prior beliefs.

So far, we have provided a micro-foundation for mislearning from endogenous outcomes, while maintaining rational Bayesian updating in the signal processing part of the expectations formation process. We now allow for biases in *both* steps of the expectation formation process, and distinguish between biases that arise from learning from (endogenous) prices, and those that arise from learning from (exogenous) fundamentals. Examples of the former type of bias include models of naïve inference (Glaeser and Nathanson 2017), PET, and K -level thinking in inference. Examples of the latter include models of over-confidence (Daniel et al. 2001), dismissiveness (Banerjee 2011), inattention (Gabaix 2019), and diagnostic expectations (Bordalo et al. 2019).

Studying these biases within a general framework serves two different purposes. First, it makes it possible to nest other biases commonly used in the behavioral economics literature, thus allowing us to clarify how the insights on instability and inelastic demands we have uncovered so far differ across biases. Second, it allows us to study how these two classes of biases might interact in a very natural way.

3.1 Biases in Learning from (Exogenous) Fundamentals

Biases that arise from learning from (exogenous) fundamentals can effectively be thought of as biases in Bayesian updating: given a fundamental signal, agents put the wrong weight on that signal in forming their posterior beliefs. In what follows, we show that many biases in learning from (exogenous) fundamentals can be expressed as follows:

$$\mathbb{E}[v|\mathcal{I}_i] = \hat{g}s_i + (1 - \hat{g})\mu_0 \tag{30}$$

$$\mathbb{V}[v|\mathcal{I}_i] = (\hat{\tau}_s + \tau_0)^{-1} \tag{31}$$

where $s_i = s$ for $i = I$, and $s_i = \tilde{s}$ for $i = U$, $\hat{\tau}_s$ is traders' perception of the precision of their signal, and \hat{g} is the weight that traders put on their signal in updating their posterior beliefs. When traders do (rational) Bayesian updating, then $\hat{\tau}_s = \tau_s$ and $\hat{g} = g \equiv \frac{\tau_s}{\tau_s + \tau_0}$. Instead, $\hat{g} > g$ and $\hat{\tau}_s > \tau_s$ push traders to over-react to their signal, and conversely $\hat{g} < g$ and $\hat{\tau}_s < \tau_s$ push traders to under-react to their signal. When computing agents' demand functions (conditional on traders' signals) and solving for market clearing, these biases affect how strongly traders' signals get incorporated into prices in the true price function, leading to $\hat{\alpha} \neq \alpha$ and $\hat{\beta} \neq \beta$. For ease of illustration, in what follows we let $Z = 0$, and we solve for the general case in Online Appendix B.2.

Definition 2 (Mislearning from (Exogenous) Fundamentals). *Let α and β be the coefficients that determine the influence on prices of informed and uninformed traders' signals when they perform (rational) Bayesian updating. When traders have biases in learning information from exogenous fundamentals, the true price function can be expressed as:*

$$P^{True}(s, \tilde{s}) = \hat{\alpha}s + \hat{\beta}\tilde{s} + (1 - \hat{\alpha} - \hat{\beta})\mu_0, \quad (32)$$

where $\hat{\alpha} \neq \alpha$, and $\hat{\beta} \neq \beta$ capture behavioral departures from the rational benchmark.

3.1.1 Nesting Biases in Learning from (Exogenous) Fundamentals

While the formulation in Definition 2 may seem restrictive, it nests a large number of biases often used in the literature. We briefly show below how one can fit over-confidence (Daniel et al. 2001), dismissiveness (Banerjee 2011), diagnostic expectations (Bordalo et al. 2019), and inattention (Gabaix 2019) using the setup and notation of Section 2.

Example 1 (Over-Confidence). *Agents are over-confident with respect to their private signal when they believe that its precision is $\hat{\tau}_s = \kappa^{oc}\tau_s$, and $\hat{g} = \frac{\kappa^{oc}\tau_s}{\kappa^{oc}\tau_s + \tau_0}$ with $\kappa^{oc} > 1$. Assuming, as in (Daniel et al. 2001), that agents use the correct precision for the signals they extract from prices, the resulting coefficients are:*

$$\hat{\alpha}^{oc} = \frac{\phi\kappa^{oc}\tau_s}{\tau_0 + (\phi\kappa^{oc} + (1 - \phi))\tau_s} \quad ; \quad \hat{\beta}^{oc} = \frac{(1 - \phi)\tau_s}{\tau_0 + (\phi\kappa^{oc} + (1 - \phi))\tau_s} \quad (33)$$

leading to $\hat{\alpha}^{oc} > \alpha$ and $\hat{\beta}^{oc} < \beta$.

Example 2 (Dismissiveness). Agents are dismissive when they believe that the precision of extracted signals is $\hat{\tau}_s = \kappa^{dis}\tau_s$ and $\hat{g} = \frac{\kappa^{dis}\tau_s}{\kappa^{dis}\tau_s + \tau_0}$, with $\kappa^{dis} < 1$. The resulting coefficients are:

$$\hat{\alpha}^{dis} = \frac{\phi\tau_s}{\tau_0 + (\phi + (1 - \phi)\kappa^{dis})\tau_s} \quad ; \quad \hat{\beta}^{dis} = \frac{(1 - \phi)\kappa^{dis}\tau_s}{\tau_0 + (\phi + (1 - \phi)\kappa^{dis})\tau_s} \quad (34)$$

leading to $\hat{\alpha}^{dis} > \alpha$ and $\hat{\beta}^{dis} < \beta$.

Example 3 (Diagnostic Expectations). Agents are diagnostic when they put too much weight on recent news, which in our framework means that all agents use $\hat{\tau}_s = \tau_s$ and $\hat{g} = \theta g$, with $\theta > 1$ capturing the degree of diagnosticity. The resulting coefficients are:

$$\hat{\alpha}^{de} = \frac{\phi\theta\tau_s}{\tau_0 + \tau_s} \quad ; \quad \hat{\beta}^{de} = \frac{(1 - \phi)\theta\tau_s}{\tau_0 + \tau_s} \quad (35)$$

leading to $\hat{\alpha}^{de} > \alpha$ and $\hat{\beta}^{de} > \beta$.

Example 4 (Inattention). Agents are inattentive in the sense of [Gabaix \(2019\)](#) when they update their beliefs only partially towards the direction of the signal: $\hat{\tau}_s = m\tau_s$ and $\hat{g} = \frac{m\tau_s}{m\tau_s + \tau_0}$, with $m < 1$ capturing the degree of inattention. The resulting coefficients are:

$$\hat{\alpha}^{in} = \frac{\phi m\tau_s}{\tau_0 + m\tau_s} \quad ; \quad \hat{\beta}^{in} = \frac{(1 - \phi)m\tau_s}{\tau_0 + m\tau_s} \quad (36)$$

leading to $\hat{\alpha}^{in} < \alpha$ and $\hat{\beta}^{in} < \beta$.

Furthermore, if one assumes that all agents correctly perform the inference problem, such that $\tilde{s} = s$ as in [Daniel et al. \(2001\)](#), [Banerjee \(2011\)](#) or [Bordalo et al. \(2020\)](#), the equilibrium price function is simply given by:

$$P(s) = (\hat{\alpha} + \hat{\beta})s + (1 - \hat{\alpha} - \hat{\beta})\mu_0. \quad (37)$$

Since $\hat{\alpha} + \hat{\beta} < 1$ is bounded above by 1, (37) makes clear that, in the absence of misinference, biases in learning from exogenous signals can *not* give rise to the type of instability that we highlighted in Section 2.3 with $\partial P/\partial s \rightarrow \infty$, and the wedges between the REE and behavioral outcomes are bounded for a given κ^{oc} , κ^{dis} , θ , or m .

Moreover, we see that whenever $\hat{\alpha} + \hat{\beta} > \alpha + \beta$ equilibrium outcomes are too sensitive to new information, leading to over-reaction, as in models of over-confidence, or diagnostic expectations. Conversely, when $\hat{\alpha} + \hat{\beta} < \alpha + \beta$ equilibrium outcomes under-react to new information, as in models of dismissiveness or inattention.²⁸

3.2 Biases in Learning from (Endogenous) Prices

Biases in inference manifest themselves in the model of the world agents use to extract information from prices, as expressed in the next definition.²⁹

Definition 3 (Mislearning from (Endogenous) Prices). *Let γ be the sensitivity of the equilibrium price to new information when traders infer the right information from prices. We capture biases in learning information from prices with the following (linear) mapping:*

$$P^{Mis}(\tilde{s}) = \tilde{\gamma}\tilde{s} + (1 - \tilde{\gamma})\mu_0 \quad (38)$$

where $\tilde{\gamma} \neq \gamma$ captures deviations from correct inference. Inverting this mapping yields the following extracted signal: $\tilde{s} = \frac{1}{\tilde{\gamma}}P - \left(\frac{1-\tilde{\gamma}}{\tilde{\gamma}}\right)\mu_0$.

If traders are only biased in the way they learn information from (endogenous) outcomes, such that $\hat{\alpha} = \alpha$ and $\hat{\beta} = \beta$, then $\tilde{\gamma} < \alpha + \beta$ leads to over-reaction, while if $\tilde{\gamma} > \alpha + \beta$ leads to under-reaction. Intuitively, if uninformed traders think that prices are less(more) responsive to news than they really are, they attribute a given price change to more(less) extreme news. Their posterior beliefs are then given by:

$$\mathbb{E}[v|\mathcal{I}_i] = g \left(\frac{1}{\tilde{\gamma}}P - \left(\frac{1-\tilde{\gamma}}{\tilde{\gamma}}\right)\mu_0 \right) + (1-g)\mu_0 \quad (39)$$

$$\mathbb{V}[v|\mathcal{I}_i] = (\tau_s + \tau_0)^{-1} \quad (40)$$

²⁸Relatedly, cursed traders in [Eyster et al. \(2019\)](#) and [Mondria et al. \(2022\)](#) infer signals from prices, but process the signal as if they thought it were noisier than it is. In our framework, this is nested as a bias in learning from fundamentals: agents behave as if they were able to infer an unbiased signal s from prices, but they underweight this signal in forming posterior beliefs.

²⁹Our framework also nests the type of behavior modeled in [Haddad et al. \(2021\)](#), where investors' demand elasticity depends on the aggregate demand elasticity. For example, with partial equilibrium thinking, uninformed traders' demand elasticity depends on their perception of the aggregate elasticity.

3.2.1 Nesting Biases in Learning from (Endogenous) Prices

Biases in inference, as PET in Section 2, manifest themselves in the $\tilde{\gamma}$ coefficient. Examples of misinference include K-Level Thinking, or using the incorrect mapping to infer information due to mistaken beliefs about the parameters of the model, such as ϕ and τ_s .

Example 5 (K-Level Thinking). *K-Level thinking agents believe that all other U agents are (K - 1)-level thinkers, and that the price they observe is generated by the (K - 1)-level thinking equilibrium price function:*

$$P^{Mis}(\tilde{s}) = \tilde{\gamma}_{K-1}\tilde{s} + (1 - \tilde{\gamma}_{K-1})\mu_0 \quad (41)$$

where $\tilde{\gamma}_{K-1} = \frac{\tau_s}{\tau_s + \tau_0 - (-\frac{1-\phi}{\phi})^K \tau_0}$ is the sensitivity of the price to the true signal in the (K - 1)-level thinking equilibrium. Notice that if $\phi < 0.5$, the bias in the mapping used by uninformed traders is increasing in the level of reasoning K , $\partial(\tilde{\gamma}_{K-1} - \gamma^{REE})/\partial K > 0$. This example illustrates that having higher levels of rational higher order beliefs does not guarantee higher welfare, and may in fact be detrimental (Bohren and Hauser 2021).³⁰

Example 6 (Misunderstanding ϕ). Let $\tilde{\phi}$ be PET agents' perception of the fraction of informed agents. PET traders then believe that the price they observe is generated by:

$$P_{\tilde{\phi}}^{Mis}(\tilde{s}) = \frac{\tilde{\phi}\tau_s}{\tilde{\phi}\tau_s + \tau_0}\tilde{s} + \frac{\tau_0}{\tilde{\phi}\tau_s + \tau_0}(\mu_0 - AZ\tau_0^{-1}) \quad (42)$$

When $0 < \tilde{\phi} < \phi$, PET agents believe that the price is even less sensitive than in our benchmark case, amplifying the results. The opposite is true for $\tilde{\phi} > \phi$. Moreover, notice that $\tilde{\phi} = 0$ is equivalent to the case where uninformed traders are fully cursed (they think prices do not contain any relevant information), and $\tilde{\phi} = 1$ recovers the REE mapping.

Example 7 (Misunderstanding τ_s). Let $\tilde{\tau}_s$ be PET agents' perception of the precision of the signal received by informed agents. PET traders then believe that the price they observe is generated by:

$$P_{\tilde{\phi}}^{Mis}(\tilde{s}) = \frac{\phi\tilde{\tau}_s}{\phi\tilde{\tau}_s + \tau_0}\tilde{s} + \frac{\tau_0}{\phi\tilde{\tau}_s + \tau_0}(\mu_0 - AZ\tau_0^{-1}) \quad (43)$$

³⁰See Online Appendix D for a detailed and complete analysis of K-level thinking.

When $\tilde{\tau}_s < \tau_s$, PET agents believe that the price is even less sensitive than in our benchmark case, amplifying the misinference. The opposite is true for $\tilde{\tau}_s > \tau_s$.

3.3 Implications for Demand Elasticities

Before studying the interaction of mislearning from prices and mislearning from fundamentals, it is worth highlighting that these two classes of biases have different implications for demand elasticities.³¹ To see why this is the case, we re-write the slope of agents' demand curves as follows:

$$-\frac{\partial X_i}{\partial P} = \frac{1}{A\mathbb{V}[v|\mathcal{I}_i]} \left(1 - \frac{\partial \mathbb{E}[v|\mathcal{I}_i]}{\partial P} \right) \quad (44)$$

where $\frac{\partial \mathbb{E}[v|\mathcal{I}_i]}{\partial P} = 0$ if traders do not learn from prices, and $\frac{\partial \mathbb{E}[v|\mathcal{I}_i]}{\partial P} > 0$ if instead they do.

Starting with the case with no learning from prices, (44) shows that biases in learning information from (exogenous) fundamentals can only affect demand elasticities by changing agents' confidence levels and risk-bearing capacities. For example, a bias like diagnostic expectations only affects the first moment of agents' beliefs, therefore leaving the slope of agents' demand curves unchanged. Alternatively, over-confident agents perceive lower uncertainty than they would under the rational benchmark ($\mathbb{V}[v|\mathcal{I}_i] < \mathbb{V}^{REE}[v|\mathcal{I}_i]$), and therefore *trade more aggressively against any perceived mispricing*. However, this form of over-reaction leads to more *elastic* demand functions. Overconfidence implies a positive relationship between the extent of excess volatility and how steep demand curves are, making the growing evidence of inelastic markets even more puzzling.

Turning to the case with learning from prices (and rational Bayesian updating), (44) shows that in this case the informational role of prices leads to more inelastic demand curves than without learning from prices. Intuitively, following a price rise, uninformed traders become more optimistic, and this *dampens* their desire to decrease their holdings due to the asset becoming more expensive. In the rational expectations equilibrium these two forces perfectly cancel out ($\frac{\partial \mathbb{E}[v|\mathcal{I}_i]}{\partial P} = 1$), leading to *perfectly inelastic* demand curves. If instead traders mistakenly think that prices are less responsive to news than under the

³¹In the following discussion we focus on models that deliver over-reaction, as models of under-reaction are at odds with the other asset pricing puzzles we explored in Section 2.5.

rational mapping ($\frac{\partial \mathbb{E}[v|\mathcal{I}_i]}{\partial P} > 1$), the informational role of prices is stronger than the scarcity role, leading to *upward sloping* demand curves and over-reaction, as with partial-equilibrium thinking.

This analysis highlights how mislearning from prices and mislearning from fundamentals have different implications for the role of over-reaction in contributing to more inelastic markets. While over-reaction from mislearning from prices contributes to more inelastic markets, over-reaction from mislearning from fundamentals alone either predicts the opposite relationship, or no relationship at all. Online Appendix F shows that mislearning from prices contributes to more inelastic markets even when prices are partially revealing: if traders are unable to distinguish between fundamental and non-fundamental price changes, they respond to both in the same way. When this is the case, learning from prices can translate into more inelastic changes in holdings even in response to non-fundamental price changes, which are mistakenly attributed to fundamental news (Mendel and Shleifer 2012).

3.3.1 Back of the Envelope Calibration

To relate our results to empirical estimates of demand elasticities in the literature, and to understand how strong our channel of interest might be in contributing to more inelastic markets, it is useful to express the slope of agent i 's demand curve in percentage terms:

$$-\frac{\partial \log X_i}{\partial \log P} = 1 - \frac{\partial \log \mathbb{E}[r|\mathcal{I}_i]}{\partial \log P} = 1 + \left(\frac{1 + \mathbb{E}[r|\mathcal{I}_i]}{\mathbb{E}[r|\mathcal{I}_i]} \right) \left(1 - \frac{\partial \log \mathbb{E}[v|\mathcal{I}_i]}{\partial \log P} \right) \quad (45)$$

where $\mathbb{E}[r|\mathcal{I}_i] \equiv \frac{\mathbb{E}[v|\mathcal{I}_i] - P}{P}$ is the expected excess return of the risky asset, and where for simplicity, we assume that $\frac{\partial \log AV[r|\mathcal{I}_i]}{\partial \log P} \approx 0$.

If we assume that the equity premium is around 4.4%, this makes clear that with no learning from prices ($\frac{\partial \log \mathbb{E}[v|\mathcal{I}_i]}{\partial \log P} = 0$), the macro elasticity of demand is simply $-\frac{\partial \log X_i}{\partial \log P} = 1 + \frac{1.044}{0.044} = 25$, which is consistent with previous theory implied estimates. As extensively reviewed in Gabaix and Koijen (2022), this is two orders of magnitude higher than empirical estimates, which find a macro elasticity of around 0.2.

However, notice how this calibration with no learning from prices also implies that the pass-through of prices to expected returns is approximately one ($\frac{\partial \mathbb{E}[r|\mathcal{I}_i]}{\partial \log P} = -1.044$). As also

noted by [Davis et al. \(2023\)](#), if prices do not *cause* changes in agents’ expected future payoffs, then an exogenous 1% increase in prices must lead to a 1% decrease in expected returns. If instead price changes do *causally* influence investors’ expectations of next period payoff, then there is an additional term in (45) that ensures that the pass-through of prices to expected returns is less than one-for-one, and this contributes to more inelastic markets. A point worth highlighting is that the direction of causality from prices to expected returns here is important, as changes in expected future payoffs that are not *caused* by price changes would show up in the level of the demand curve, and not in its slope.³²

Empirically, [Davis et al. \(2023\)](#) show that the pass-through of prices to expected returns is much lower than one, which implies that price changes do influence expectations of future payoffs. Specifically, they show that at the micro level the pass-through of prices to expected returns is around -0.06 , and this contributes to lowering the theory implied micro elasticity of demand by an order of magnitude. At the macro level, [De La O and Myers \(2021\)](#) use survey data and find estimates of $\frac{\partial \mathbb{E}[r]}{\partial \log P} \approx 0.05$, while [Nagel and Xu \(2023\)](#) find that even when $\frac{\partial \mathbb{E}[r]}{\partial \log P}$ is negative, it is quantitatively small. According to our model, this would translate into a macro elasticity of demand of around 1, which is also an order of magnitude closer to the empirical estimates of 0.2. This suggests that the channel we are studying could be quantitatively meaningful. Of course, these estimates of the elasticity of expected returns do not control for the direction of causality required in our calibration. [Chaudhry \(2022\)](#) provides more direct empirical support for learning from prices by introducing subjective beliefs data in demand based asset pricing, and finding that exogenous price increases do cause analysts to update their beliefs of expected future payoffs, as price increases lead to upward revisions in both their short term earning per share expectations and their long term earning growth expectations.³³ A more detailed analysis of the causal influence of prices on expected returns would help in understanding the role of learning from prices in contributing to more inelastic markets, therefore shedding further light on both beliefs and holdings.

³²For instance, [Greenwood and Hanson \(2015\)](#) present a model of competition neglect where investors neglect other investors’ responses to a fundamental shock. This leads to an excessive change in investors’ expected future payoffs that does not depend on the current price. Therefore, this leads to an excessive *shift* in investors’ demand curves, while leaving the elasticity of demand unchanged.

³³Importantly, our paper shows that the causal effect of prices on beliefs does not negate the effect of biased beliefs on prices. Instead, mislearning from prices may contribute to more inelastic markets, further amplifying other forms of mispricing.

3.4 Equilibrium and Interaction

As before, the unique equilibrium is given by the intersection of the two mappings in (32) and (38). This yields the following expressions for the equilibrium price and extracted signal:

$$P = \frac{\hat{\alpha}}{1 - \frac{\hat{\beta}}{\tilde{\gamma}}} s + \left(1 - \frac{\hat{\alpha}}{1 - \frac{\hat{\beta}}{\tilde{\gamma}}}\right) \mu_0 \quad (46)$$

$$\tilde{s} = s + \left(\frac{\hat{\alpha} + \hat{\beta} - \tilde{\gamma}}{\tilde{\gamma} - \hat{\beta}}\right) (s - \mu_0) \quad (47)$$

We start by generalizing the conditions for stability uncovered in PET, as we are interested in studying the properties of stable equilibria.³⁴ The relevant parameters in determining the strength of the feedback effect are $\tilde{\gamma}$ and $\hat{\beta}$. Instead, stability is not directly related to the rational parameters α and β .

Proposition 3 (Stability with General Model Misspecification). *Given a general linear mapping, as in (38) and a true model of the world as in (32), the equilibrium is stable if and only if $\frac{\hat{\beta}}{\tilde{\gamma}} < 1$.*

The expression in (47) makes clear that for agents to perform correct inference we must have $\tilde{\gamma} = \hat{\alpha} + \hat{\beta}$: only then do U agents use the right mapping to infer information from prices, allowing them to recover the correct signal $\tilde{s} = s$, for any true signal s . Instead, when $\tilde{\gamma} \neq \hat{\alpha} + \hat{\beta}$, agents' signal extraction results in misinference. Specifically, when $\tilde{\gamma} < \hat{\alpha} + \hat{\beta}$, U agents think that the price is less sensitive to new information than in reality, and therefore they extract a signal which is more extreme, as in PET, and conversely for $\tilde{\gamma} > \hat{\alpha} + \hat{\beta}$. Online Appendix E shows how this point generalizes to a model with n types of heterogeneous agents who all use a different (linear) mapping to infer information from prices.

Proposition 4 (Correct Inference from Prices with Biases in Learning from Exogenous Fundamentals). *When $\hat{\alpha} \neq \alpha$, uninformed traders extract the right information from prices if and only if $\tilde{\gamma} = \hat{\alpha} + \hat{\beta}$. If $\tilde{\gamma} < \hat{\alpha} + \hat{\beta}$, traders infer a signal that is more extreme than the true one. If $\tilde{\gamma} > \hat{\alpha} + \hat{\beta}$, traders infer a signal that is less extreme than the true one.*

³⁴When the equilibrium is unstable, all comparative statics are reversed. This is simply a manifestation of Samuelson's Correspondence Principle (Samuelson 1947).

The equilibrium effects of the interaction between biases that arise from processing fundamental signals, and those that arise from learning from prices are shown in Figure 2, where the two thick dashed lines correspond to regions where agents do rational Bayesian updating ($\hat{\alpha} = \alpha$, $\hat{\beta} = \beta$) and correct inference from prices ($\tilde{\gamma} = \hat{\alpha} + \hat{\beta}$), respectively.³⁵ The REE then corresponds to the intersection of these two lines. Away from this intersection, the two lines split the plane into four quadrants, where the two departures from rationality interact in shaping the properties of equilibrium outcomes: biases such that $\hat{\alpha} + \hat{\beta} > \alpha + \beta$ and $\tilde{\gamma} < \hat{\alpha} + \hat{\beta}$ both push towards over-reaction, while reversing these inequalities pushes towards under-reaction. As is clear from Figure 2, these forces may either point in the same direction, in which case they compound and amplify each other (quadrants 1 and 3), or they push in opposite directions (quadrants 2 and 4). Proposition 5 shows that the resulting properties of equilibrium outcomes depend on the relative strength of the two mistakes.

Proposition 5 (Over- and Under-reaction with General Model Misspecification). *Given a general linear mapping as in (38), a true model of the world as in (32), and rational coefficients α and β , a stable equilibrium displays over-reaction if:*

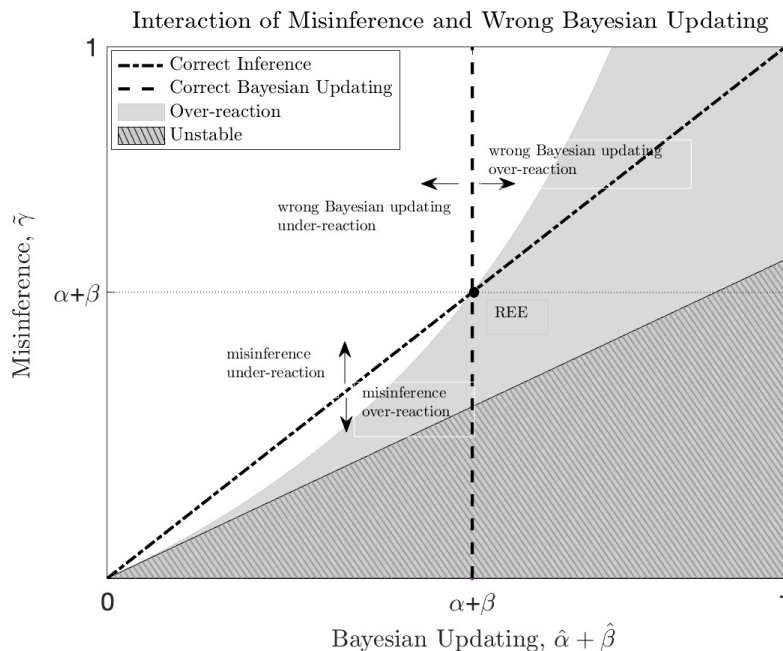
$$\frac{\hat{\alpha}}{1 - \frac{\hat{\beta}}{\tilde{\gamma}}} > \alpha + \beta \iff \tilde{\gamma} < (\alpha + \beta) \left(\frac{\hat{\beta}}{\beta - (\hat{\alpha} - \alpha)} \right) \quad (48)$$

and under-reaction if the inequality is reversed. When $\hat{\alpha} = \alpha$ and $\hat{\beta} = \beta$, this inequality reduces to: $\tilde{\gamma} < \alpha + \beta = \gamma_{REE}$.

A key point to notice is just how plausible it is to have both mistakes at play in any setup where agents mislearn from fundamentals *and* also extract information from endogenous outcomes. Such models generally assume that agents do correct inference, implicitly assuming that agents are aware of the mistakes of others. For example, a common assumption in the over-confidence literature is that agents perform correct signal extraction, and must therefore be aware that others are overconfident, even though they fail to recognize that they are themselves overconfident regarding their own information (e.g. Odean 1998, or Daniel et al. 1998). If instead agents believe that they live in a rational world, thus failing to

³⁵Details on how we vary $\hat{\alpha}$ and $\hat{\beta}$ are illustrated in Online Appendix B.1.

Figure 2: General Model Misspecification: mislearning from prices, mislearning from fundamentals, and their interaction. This figure illustrates how mislearning from fundamentals and mislearning from prices interact in generating under- and over-reaction. The dashed vertical line at $\hat{\alpha} + \hat{\beta} = \alpha + \beta$ corresponds to the region where all agents perform rational Bayesian updating, and the dash-dotted diagonal line with $\tilde{\gamma} = \hat{\alpha} + \hat{\beta}$ corresponds to the region where agents perform correct inference from prices. The REE is at the intersection of these two lines. The horizontal axis captures mislearning from fundamentals, while the vertical axis captures mislearning from prices. The horizontal dotted line at $\tilde{\gamma} = \alpha + \beta$ corresponds to the case where agents infer information from prices under the (possibly incorrect) assumption that all agents are rational. The white and the gray areas illustrate combinations of parameters which lead to under-reaction and over-reaction, respectively. The dark gray area shows regions where the equilibrium is unstable.



realize that all other agents are also overconfident, they underestimate the sensitivity of the equilibrium price to new information, and they extract a more extreme signal than the true one, leading to even stronger departures from the REE. Similarly, in [Bordalo et al. \(2020\)](#), diagnostic agents infer the correct signal from prices, which requires agents to recognize that other agents are diagnostic. In what follows, we show that misinference can further amplify biases in Bayesian updating and even result in extreme outcomes and unstable regions.

Specifically, Figure 2 shows more formally how discrepancies from the REE arising from wrong Bayesian updating are amplified when agents misinfer information from prices because they assume the world is rational when it is not.³⁶ The thin dotted line with $\tilde{\gamma} = \alpha + \beta$ in Figure 2 corresponds to the region where agents infer information from prices under the

³⁶Relatedly, [Bohren and Hauser \(2021\)](#) show in a general social learning setup that failing to account for others' biases can have an equally severe impact on learning as having the bias itself.

(possibly incorrect) assumption that the world is rational. Importantly, this line always lies either in the top-left, or in the bottom-right quadrants described above, and wrong Bayesian updating and misinference amplify each other in both these regions. Figure 2 also shows that this type of misinference (with $\tilde{\gamma} = \alpha + \beta \neq \hat{\alpha} + \hat{\beta}$) can even lead to unstable outcomes, as is clear from the fact that, for large enough $\hat{\alpha} + \hat{\beta}$ the thin dotted line enters the unstable region. Online Appendix B.1 explores an example in further detail.

Corollary 3 (Misinference from (mistakenly) assuming rational learning from fundamentals). *Consider a scenario where agents have biases in learning from fundamentals, $\hat{\alpha} + \hat{\beta} \neq \alpha + \beta$. Misinference due to the incorrect assumption that the world is rational, $\tilde{\gamma} = \alpha + \beta$, amplifies the bias due to mislearning from fundamentals:*

- *If $\hat{\alpha} + \hat{\beta} < \alpha + \beta$, $\tilde{\gamma} = \alpha + \beta$ amplifies under-reaction relative to $\tilde{\gamma} = \hat{\alpha} + \hat{\beta}$;*
- *If $\hat{\alpha} + \hat{\beta} > \alpha + \beta$, $\tilde{\gamma} = \alpha + \beta$ either amplifies over-reaction relative to $\tilde{\gamma} = \hat{\alpha} + \hat{\beta}$, or it makes the equilibrium unstable.*

This discussion bears important quantitative implications for models which aim to assess the importance of these behavioral biases in empirical settings: smaller departures from rational Bayesian updating may be needed to obtain empirically meaningful departures from rationality. Similarly, we should expect the same bias in learning from fundamentals to have a greater influence on aggregate outcomes in settings where agents infer information from endogenous outcomes. Intuitively, by making demands more inelastic, biases in learning information from prices amplify other types of shocks, including shifts in the aggregate demand function due to biases in learning from fundamentals.

Finally, unpacking the true driver of these departures from rationality is important as the two types of biases introduce different frictions, and have different policy implications: while misinference can be attenuated with greater transparency and better communication, as in [Angeletos and Sastry \(2021\)](#), errors in Bayesian updating would be unaffected by these measures and generally require more direct changes in agents' incentives with, for instance, taxes and subsidies as in [Farhi and Gabaix \(2020\)](#).

4 Extensions

4.1 Heterogenous Agents

So far, we have considered the case where all uninformed agents extract information from prices using the same misspecified model of the world. [Giglio et al. \(2021\)](#) document a large amount of heterogeneity among agents' actions and beliefs. The tractability of our model allows us to consider the interaction of heterogeneous agents, who differ in their level of sophistication in extracting information from prices. The key point in solving these models is to notice that we have to specify the mapping used by each type of agent who extracts information from prices. Online Appendix E develops a setup and allows for N types of agents, each with their own different (linear) misspecified model.

When we extend our setup in this way, we find that the main insights in the paper generalize to a setup with heterogeneous agents. At the *individual* level, agents who believe the equilibrium price is less(more) responsive than it *truly* is attribute a given price change to a more(less) extreme signal, and overreact(underreact).³⁷ This also corroborates the results in Section 3 that agents who believe the world is rational need not necessarily be better off. At the *aggregate* level, outcomes display over-(under-)reaction if, on average, agents believe the equilibrium price is less(more) sensitive than in the *REE*.³⁸ Therefore, aggregate level rationality need not imply individual level rationality if agents' mistakes cancel out on average. Finally, we show that with heterogeneous levels of thinking, having a higher level of rationality of higher order beliefs does not guarantee a higher level of realized welfare.

³⁷Therefore, when there are heterogeneous types of uninformed agents, to understand responses at the individual level, one needs to compare agents' misspecified mappings with the resulting equilibrium mapping, and not with the REE one. For example, if the equilibrium exhibits over-reaction, an agent who believes the world is rational overreacts. Similarly, even agents who think that there is aggregate level over-reaction may overreact if they underestimate the extent of aggregate level over-reaction. This is reminiscent of the results in Section 3.

³⁸In Online Appendix E, we maintain rational Bayesian updating for simplicity. We could easily generalize these results further by allowing for non-rational Bayesian updating, as in Section 3. With non-rational Bayesian updating, correct inference is not enough to ensure the recovery of the REE at the aggregate level.

4.2 Partially Revealing Prices and Symmetric Private Signals

In Online Appendix F we relax the assumption that prices are fully revealing, and allow for the supply of the risky asset to be stochastic, as in [Diamond and Verrecchia \(1981\)](#). In Online Appendix G we maintain the assumption of partially revealing prices and consider the symmetric case where all agents receive a noisy private signal of the risky asset’s payoff. The main intuitions pertaining to partial equilibrium thinking go through, thus showing that our results do not necessarily rely on the particular information structure that we adopt in the baseline model.

5 Conclusions

We study the implications of biases arising from mislearning from equilibrium prices, and contrast them to biases arising from mislearning from exogenous fundamentals. To do so, we break down the expectation formation process in two steps: a signal extraction step, where traders learn fundamental information from prices, and a signal processing step, where traders compute their posterior beliefs by combining their inferred information with their prior. This allows us to study the different implications of mislearning from prices and mislearning from fundamentals, and to then turn to their interaction.

We start by providing a micro-foundation for mislearning from prices while maintaining rationality in the second step of the expectation formation process. Specifically, building on experimental evidence that traders under-estimate the extent of other agents’ inference from aggregate outcomes, as well as growing empirical evidence in finance suggesting that fundamental and return expectations may be interlinked, we develop a theory of partial equilibrium thinking, where traders learn fundamental information from prices, but fail to realize that all other traders do so too. We show that PET provides a micro-foundation for over-reaction to news, while also generating endogenously heterogeneous beliefs. These two features allow us to speak to the asset pricing puzzles of excess volatility, excess volume, and return predictability. Additionally, we show that PET leads to upwards sloping demand curves for uninformed traders, therefore contributing to more inelastic markets. The degree of over-reaction and inelasticity then depend on the composition of traders in the market,

and on the informativeness of new information.

Turning to the distinction between mislearning from prices and mislearning from fundamentals, there are two key points worth highlighting. First, mislearning from prices embeds a very natural two-way feedback between prices and beliefs, which can lead to arbitrarily large deviations from rationality, even for a fixed bias and bounded signal. In the limit when the feedback effect becomes explosive, the equilibrium is unstable. Second, mislearning from prices and mislearning from fundamentals have very different implications when it comes to understanding the growing evidence on inelastic markets. Unlike mislearning from fundamentals, mislearning from prices can generate both over-reaction and more inelastic markets relative to the rational expectations benchmark.

Finally, we study the interaction of mislearning from prices and mislearning from fundamentals. These two classes of biases are in no way mutually exclusive. Instead, they interact in very natural ways, and mislearning from prices can vastly amplify biases from mislearning from fundamentals. This has important quantitative implications as it suggests that biases in learning from prices can help match empirical evidence at odds with rational expectations without relying on extreme forms of biases in signal processing.

Acknowledgments: Previous versions have been circulated under the title “Partial Equilibrium Thinking in General Equilibrium.” We thank Thomas Chaney (the Editor), four anonymous referees, Nicholas Barberis, Aislinn Bohren (discussant), John Campbell, Kent Daniel, Ester Faia, Emmanuel Farhi, Xavier Gabaix, Tristan Gagnon-Bartsch, Nicola Genaioli, Niels Gormsen, Thomas Graeber, Robin Greenwood, Sam Hanson, Oliver Hart, Ryota Iijima, Ralph Koijen, David Laibson, Eben Lazarus, Ian Martin, Matthew Rabin, Josh Schwartzstein, Andrei Shleifer, Jeremy Stein, Tomasz Strzalecki, Adi Sunderam, Dimitri Vayanos, Luis Viceira, Jessica Wachter and seminar participants at Harvard (Finance, Theory, Behavioral, Contracts and Organizations Lunches), LSE, UCLA Anderson, Chicago Fed, Berkeley Haas, Duke Fuqua, TSE; as well as at the 10th Miami Behavioral Finance conference, 12th World Congress of the Econometric Association, NBER (Behavioral Finance), SITE (Psychology and Economics), Barcelona Summer Forum (Preferences and Bounded Rationality), briq Beliefs Workshop, and ESSFM (Asset Pricing) for useful com-

ments. Francesca Bastianello is grateful for support from the Alfred P. Sloan Foundation Pre-doctoral Fellowship in Behavioral Macroeconomics, awarded through the NBER.

Data Availability Statement. The code that generates all figures in the paper is available on Zenodo at: <https://doi.org/10.5281/zenodo.10780393>.

References

- Amador, M. and Weill, P.-O. (2010), ‘Learning from prices: Public communication and welfare’, *Journal of Political Economy* **118**(5), 866–907.
- Angeletos, G.-M. and Lian, C. (2017), Dampening general equilibrium: From micro to macro, Technical report, National Bureau of Economic Research.
- Angeletos, G.-M. and Sastry, K. A. (2021), ‘Managing expectations: Instruments versus targets’, *The Quarterly Journal of Economics* **136**(4), 2467–2532.
- Angeletos, G.-M. and Werning, I. (2006), ‘Crises and prices: Information aggregation, multiplicity, and volatility’, *American Economic Review* **96**(5), 1720–1736.
- Augenblick, N. and Lazarus, E. (2023), ‘A new test of excess movement in asset prices’.
- Bae, K.-H., Stulz, R. M. and Tan, H. (2008), ‘Do local analysts know more? a cross-country study of the performance of local analysts and foreign analysts’, *Journal of financial economics* **88**(3), 581–606.
- Banerjee, A. V. (1992), ‘A simple model of herd behavior’, *The Quarterly Journal of Economics* **107**(3), 797–817.
- Banerjee, S. (2011), ‘Learning from prices and the dispersion in beliefs’, *The Review of Financial Studies* **24**(9), 3025–3068.
- Barberis, N. (2018), Psychology-based models of asset prices and trading volume, in ‘Handbook of Behavioral Economics: Applications and Foundations 1’, Vol. 1, Elsevier, pp. 79–175.
- Barberis, N., Greenwood, R., Jin, L. and Shleifer, A. (2018), ‘Extrapolation and bubbles’, *Journal of Financial Economics* **129**(2), 203–227.
- Barberis, N. and Shleifer, A. (2003), ‘Style investing’, *Journal of Financial Economics* **68**(2), 161–199.
- Bastianello, F. and Fontanier, P. (2023), ‘Partial Equilibrium Thinking, Extrapolation, and Bubbles’, *Working Paper*.
- Ben-David, I., Li, J., Rossi, A. and Song, Y. (2022), ‘Ratings-driven demand and systematic price fluctuations’, *The Review of Financial Studies* **35**(6), 2790–2838.
- Beutel, J. and Weber, M. (2022), ‘Beliefs and portfolios: Causal evidence’, *Chicago Booth Research Paper* (22-08).

- Bikhchandani, S., Hirshleifer, D. and Welch, I. (1992), ‘A theory of fads, fashion, custom, and cultural change as informational cascades’, *Journal of Political Economy* **100**(5), 992–1026.
- Boehmer, E. and Kelley, E. K. (2009), ‘Institutional investors and the informational efficiency of prices’, *The Review of Financial Studies* **22**(9), 3563–3594.
- Bohren, J. A. (2016), ‘Informational herding with model misspecification’, *Journal of Economic Theory* **163**, 222–247.
- Bohren, J. A. and Hauser, D. N. (2021), ‘Learning with heterogeneous misspecified models: Characterization and robustness’, *Econometrica* **89**(6), 3025–3077.
- Bordalo, P., Gennaioli, N., Kwon, S. Y. and Shleifer, A. (2020), ‘Diagnostic bubbles’, *Journal of Financial Economics* .
- Bordalo, P., Gennaioli, N., LaPorta, R. and Shleifer, A. (2022), Belief overreaction and stock market puzzles, Technical report, Working paper.
- Bordalo, P., Gennaioli, N., Porta, R. L. and Shleifer, A. (2019), ‘Diagnostic expectations and stock returns’, *The Journal of Finance* **74**(6), 2839–2874.
- Campbell, J. Y. (2018), *Financial decisions and markets: a course in asset pricing*, Princeton University Press.
- Campbell, J. Y. and Shiller, R. J. (1988), ‘Stock prices, earnings, and expected dividends’, *The Journal of Finance* **43**(3), 661–676.
- Chahrour, R. and Gaballo, G. (2021), ‘Learning from house prices: Amplification and business fluctuations’, *The Review of Economic Studies* **88**(4), 1720–1759.
- Chandrasekhar, A. G., Larreguy, H. and Xandri, J. P. (2020), ‘Testing models of social learning on networks: Evidence from two experiments’, *Econometrica* **88**(1), 1–32.
- Chang, Y.-C., Hong, H. and Liskovich, I. (2015), ‘Regression discontinuity and the price effects of stock market indexing’, *The Review of Financial Studies* **28**(1), 212–246.
- Chaudhry, A. (2022), Do subjective growth expectations matter for asset prices?, Working paper.
- Cochrane, J. H. (2011), ‘Presidential address: Discount rates’, *The Journal of Finance* **66**(4), 1047–1108.
- Cohen, S. and Li, S. (2022), ‘Sequential cursed equilibrium’, *arXiv preprint arXiv:2212.06025*
- Costa-Gomes, M., Crawford, V. P. and Broseta, B. (2001), ‘Cognition and behavior in normal-form games: An experimental study’, *Econometrica* **69**(5), 1193–1235.
- Crawford, V. P., Costa-Gomes, M. A. and Iriberri, N. (2013), ‘Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications’, *Journal of Economic Literature* **51**(1), 5–62.
- Da, Z., Larrain, B., Sialm, C. and Tessada, J. (2018), ‘Destabilizing financial advice: Evidence from pension fund reallocations’, *The Review of Financial Studies* **31**(10), 3720–3755.

- Daniel, K. D., Hirshleifer, D. and Subrahmanyam, A. (2001), ‘Overconfidence, arbitrage, and equilibrium asset pricing’, *The Journal of Finance* **56**(3), 921–965.
- Daniel, K., Hirshleifer, D. and Subrahmanyam, A. (1998), ‘Investor psychology and security market under-and overreactions’, *The Journal of Finance* **53**(6), 1839–1885.
- Davis, C., Kargar, M. and Li, J. (2023), ‘Why is asset demand inelastic?’.
- De La O, R. and Myers, S. (2021), ‘Subjective cash flow and discount rate expectations’, *The Journal of Finance* **76**(3), 1339–1387.
- De Long, J. B., Shleifer, A., Summers, L. H. and Waldmann, R. J. (1990), ‘Positive feedback investment strategies and destabilizing rational speculation’, *The Journal of Finance* **45**(2), 379–395.
- DeMarzo, P. M., Vayanos, D. and Zwiebel, J. (2003), ‘Persuasion bias, social influence, and unidimensional opinions’, *The Quarterly Journal of Economics* **118**(3), 909–968.
- Diamond, D. W. and Verrecchia, R. E. (1981), ‘Information aggregation in a noisy rational expectations economy’, *The Journal of Financial Economics* **9**(3), 221–235.
- Enke, B. and Zimmermann, F. (2019), ‘Correlation neglect in belief formation’, *The Review of Economic Studies* **86**(1), 313–332.
- Esponda, I. and Pouzo, D. (2016), ‘Berk–nash equilibrium: A framework for modeling agents with misspecified models’, *Econometrica* **84**(3), 1093–1130.
- Esponda, I. and Vespa, E. (2014), ‘Hypothetical thinking and information extraction in the laboratory’, *American Economic Journal: Microeconomics* **6**(4), 180–202.
- Esponda, I. and Vespa, E. (2023), ‘Contingent thinking and the sure-thing principle: Revisiting classic anomalies in the laboratory’.
- Eyster, E., Galeotti, A., Kartik, N. and Rabin, M. (2014), ‘Congested observational learning’, *Games and Economic Behavior* **87**, 519–538.
- Eyster, E. and Rabin, M. (2008), ‘Naive herding’, *Working Paper, London School of Economics* .
- Eyster, E. and Rabin, M. (2010), ‘Naive herding in rich-information settings’, *American Economic Journal: Microeconomics* **2**(4), 221–43.
- Eyster, E., Rabin, M. and Vayanos, D. (2019), ‘Financial markets where traders neglect the informational content of prices’, *The Journal of Finance* **74**(1), 371–399.
- Eyster, E., Rabin, M. and Weizsäcker, G. (2018), ‘An experiment on social mislearning’, *Working paper*.
- Eyster, E. and Weizsäcker, G. (2010), ‘Correlation neglect in financial decision-making’, *DIW Berlin Discussion Paper* .
- Farhi, E. and Gabaix, X. (2020), ‘Optimal taxation with behavioral agents’, *American Economic Review* **110**(1), 298–336.

- Farhi, E. and Werning, I. (2019), ‘Monetary policy, bounded rationality, and incomplete markets’, *American Economic Review* **109**(11), 3887–3928.
- French, K. R. (2008), ‘Presidential address: The cost of active investing’, *The Journal of Finance* **63**(4), 1537–1573.
- Frick, M., Iijima, R. and Ishii, Y. (2020), ‘Misinterpreting others and the fragility of social learning’, *Econometrica* **88**(6), 2281–2328.
- Gabaix, X. (2019), Behavioral inattention, in ‘Handbook of Behavioral Economics: Applications and Foundations 1’, Vol. 2, Elsevier, pp. 261–343.
- Gabaix, X. and Koijen, R. S. (2022), ‘In search of the origins of financial fluctuations: The inelastic markets hypothesis’.
- Gabaix, X., Koijen, R. S. J., Mainardi, F., Oh, S. and Yogo, M. (2022), ‘Asset demand of u.s. households’, *Working Paper* .
- Gagnon-Bartsch, T. and Rabin, M. (2017), ‘Naive social learning, mislearning, and unlearning’, *Working Paper* .
- Gagnon-Bartsch, T., Rabin, M. and Schwartzstein, J. (2020), ‘Channeled attention and stable errors’, *Working Paper, Harvard Business School* .
- García-Schmidt, M. and Woodford, M. (2019), ‘Are low interest rates deflationary? a paradox of perfect-foresight analysis’, *American Economic Review* **109**(1), 86–120.
- Giglio, S., Maggiori, M., Stroebel, J. and Utkus, S. (2021), ‘Five facts about beliefs and portfolios’, *American Economic Review* **111**(5), 1481–1522.
- Glaeser, E. L. and Nathanson, C. G. (2017), ‘An extrapolative model of house price dynamics’, *Journal of Financial Economics* **126**(1), 147–170.
- Gompers, P. A. and Metrick, A. (2001), ‘Institutional investors and equity prices’, *The quarterly journal of Economics* **116**(1), 229–259.
- Greenwood, R. and Hanson, S. G. (2015), ‘Waves in ship prices and investment’, *The Quarterly Journal of Economics* **130**(1), 55–109.
- Greenwood, R. and Shleifer, A. (2014), ‘Expectations of returns and expected returns’, *The Review of Financial Studies* **27**(3), 714–746.
- Grossman, S. J. and Stiglitz, J. E. (1980), ‘On the impossibility of informationally efficient markets’, *The American Economic Review* **70**(3), 393–408.
- Haddad, V., Huebner, P. and Loualiche, E. (2021), How competitive is the stock market? theory, evidence from portfolios, and implications for the rise of passive investing, Working paper.
- Hartzmark, S. M. and Solomon, D. H. (2022), Predictable price pressure, Technical report, National Bureau of Economic Research.
- Hayek, F. A. (1945), ‘The use of knowledge in society’, *The American Economic Review* **35**(4), 519–530.

- Hellwig, C., Mukherji, A. and Tsyvinski, A. (2006), ‘Self-fulfilling currency crises: The role of interest rates’, *American Economic Review* **96**(5), 1769–1787.
- Hong, H., Lim, T. and Stein, J. C. (2000), ‘Bad news travels slowly: Size, analyst coverage, and the profitability of momentum strategies’, *The Journal of Finance* **55**(1), 265–295.
- Hong, H. and Stein, J. C. (1999), ‘A unified theory of underreaction, momentum trading, and overreaction in asset markets’, *The Journal of Finance* **54**(6), 2143–2184.
- Iovino, L. and Sergeyev, D. (2020), ‘Central bank balance sheet policies without rational expectations’, *CEPR Discussion Paper* .
- Jensen, T. I. (2023), ‘Subjective risk and return’, *Available at SSRN 4276760* .
- Jin, L. J. and Sui, P. (2022), ‘Asset pricing with return extrapolation’, *Journal of Financial Economics* **145**(2), 273–295.
- Keniston, D., Larsen, B. J., Li, S., Prescott, J. J., Silveira, B. S. and Yu, C. (2021), Fairness in incomplete information bargaining: Theory and widespread evidence from the field, Technical report, National Bureau of Economic Research.
- Koijen, R. S., Richmond, R. J. and Yogo, M. (2020), Which investors matter for equity valuations and expected returns?, Technical report, National Bureau of Economic Research.
- Koijen, R. S. and Yogo, M. (2019), ‘A demand system approach to asset pricing’, *Journal of Political Economy* **127**(4), 1475–1515.
- Kübler, D. and Weizsäcker, G. (2004), ‘Limited depth of reasoning and failure of cascade formation in the laboratory’, *The Review of Economic Studies* **71**(2), 425–441.
- Laarits, T. and Sammon, M. (2022), ‘The retail habitat’, *Available at SSRN 4262861* .
- Levy, G. and Razin, R. (2015), ‘Correlation neglect, voting behavior, and information aggregation’, *American Economic Review* **105**(4), 1634–45.
- Li, J. (2022), ‘What drives the size and value factors?’, *Review of Asset Pricing Studies* **12**(4), 845–885.
- Martínez-Marquina, A., Niederle, M. and Vespa, E. (2019), ‘Failures in contingent reasoning: The role of uncertainty’, *American Economic Review* **109**(10), 3437–3474.
- Mendel, B. and Shleifer, A. (2012), ‘Chasing noise’, *Journal of Financial Economics* **104**(2), 303–320.
- Mondria, J., Vives, X. and Yang, L. (2022), ‘Costly interpretation of asset prices’, *Management Science* **68**(1), 52–74.
- Moreira, A. and Muir, T. (2017), ‘Volatility-managed portfolios’, *The Journal of Finance* **72**(4), 1611–1644.
- Nagel, R. (1995), ‘Unraveling in guessing games: An experimental study’, *The American Economic Review* **85**(5), 1313–1326.
- Nagel, S. and Xu, Z. (2022), ‘Asset pricing with fading memory’, *The Review of Financial Studies* **35**(5), 2190–2245.

- Nagel, S. and Xu, Z. (2023), ‘Dynamics of subjective risk premia’, *Journal of Financial Economics* **150**(2), 103713.
- Ngangoué, M. K. and Weizsäcker, G. (2021), ‘Learning from unrealized versus realized prices’, *American Economic Journal: Microeconomics* **13**(2), 174–201.
- Niederle, M. and Vespa, E. (2023), ‘Cognitive limitations: Failures of contingent thinking’, *Annual Review of Economics* **15**, 307–328.
- Odean, T. (1998), ‘Volume, volatility, price, and profit when all traders are above average’, *The Journal of Finance* **53**(6), 1887–1934.
- Pavlova, A. and Sikorskaya, T. (2023), ‘Benchmarking intensity’, *The Review of Financial Studies* **36**(3), 859–903.
- Penczynski, S. P. (2017), ‘The nature of social learning: Experimental evidence’, *European Economic Review* **94**, 148–165.
- Peng, C. and Wang, C. (2023), ‘Factor demand and factor returns’, *Available at SSRN 3327849*.
- Petajisto, A. (2009), ‘Why do demand curves for stocks slope down?’, *Journal of Financial and Quantitative Analysis* **44**(5), 1013–1044.
- Samuelson, P. A. (1947), ‘Foundations of economic analysis.’
- Scheinkman, J. A. and Xiong, W. (2003), ‘Overconfidence and speculative bubbles’, *Journal of Political Economy* **111**(6), 1183–1220.
- Shiller, R. (1981), ‘Do stock prices move too much to be justified by subsequent changes in dividends?’, *American Economic Review* **71**, 421–736.
- Shleifer, A. (1986), ‘Do demand curves for stocks slope down?’, *The Journal of Finance* **41**(3), 579–590.
- Stahl, D. O. and Wilson, P. W. (1995), ‘On players’ models of other players: Theory and experimental evidence’, *Games and Economic Behavior* **10**(1), 218–254.
- Veldkamp, L. L. (2023), *Information choice in macroeconomics and finance*, Princeton University Press.
- Yan, X. and Zhang, Z. (2009), ‘Institutional investors and equity returns: are short-term institutions better informed?’, *The Review of Financial Studies* **22**(2), 893–924.
- Ziegelmeyer, A., March, C. and Krügel, S. (2013), ‘Do we follow others when we should? a simple test of rational expectations: comment’, *American Economic Review* **103**(6), 2633–42.

Appendices

A Proofs

To keep the notation succinct in this Section, we sometimes adopt the notation of Section 3, and replace the coefficients with the fundamental parameters when needed.

A.1 Proposition 1: Demand Elasticities with PET

When U traders learn information from prices, we can write the slope of their demand curves as:

$$\frac{\partial X_{U,K}}{\partial P} = -\frac{1}{A(\tau_s + \tau_0)^{-1}} \left(1 - \left(\frac{\tau_s}{\tau_s + \tau_0} \right) \left(\frac{1}{\gamma_K} \right) \right) \quad (\text{A.1})$$

where $K \in \{REE, PET\}$, such that $\gamma_{REE} = \frac{\tau_s}{\tau_0 + \tau_s}$ and $\gamma_{PET} = \frac{\phi\tau_s}{\phi\tau_s + \tau_0} < \frac{\tau_s}{\tau_s + \tau_0}$. It follows that $\frac{\partial X_{U,PET}}{\partial P} > 0 = \frac{\partial X_{U,REE}}{\partial P}$, so that PET leads to upward sloping demand curves for uninformed traders. The slope of informed traders' demand curve is the same under PET and under REE, $\frac{\partial X_{I,PET}}{\partial P} = -\frac{1}{A(\tau_s + \tau_0)^{-1}} = \frac{\partial X_{I,REE}}{\partial P}$. We can then write the ratio of the slopes of the aggregate demand curves under PET and under REE as:

$$\frac{\partial X_{TOT,PET}/\partial P}{\partial X_{TOT,REE}/\partial P} = 1 - (1 - \phi) \left(\frac{\frac{\tau_s}{\tau_0} + \frac{1}{\phi}}{\frac{\tau_s}{\tau_0} + 1} \right) < 1 \quad (\text{A.2})$$

which shows that PET leads to more inelastic markets than under REE, and the wedge is decreasing in ϕ and $\frac{\tau_s}{\tau_0}$. \square

A.2 Lemma 1: Stability

Using $\tilde{\gamma} = \phi\tau_s/(\phi\tau_s + \tau_0)$, we have:

$$\frac{\partial X_{TOT}}{\partial P} < 0 \iff (1 - \phi) \left(\frac{\tau_s}{\tau_s + \tau_0} \right) \frac{1}{\tilde{\gamma}} < 1 \iff \beta < \tilde{\gamma} \iff \frac{\partial P^{True}(s, \tilde{s})}{\partial \tilde{s}} < \frac{\partial P^{Mis}(\tilde{s})}{\partial \tilde{s}} \quad (\text{A.3})$$

where the last two equivalences follow from using the definitions of β and $\tilde{\gamma}$. Re-writing this inequality in terms of the primitives of the model, rearranging, and simplifying we have that:

$$\beta < \tilde{\gamma} \iff \frac{(1 - \phi)\tau_s}{\tau_s + \tau_0} < \frac{\phi\tau_s}{\phi\tau_s + \tau_0} \implies \phi^2\tau_s + 2\phi\tau_0 - \tau_0 > 0 \quad (\text{A.4})$$

$$\iff \tau_s + \tau_0 - \left(\frac{1 - \phi}{\phi} \right)^2 \tau_0 > 0 \iff \phi > \left(1 + \sqrt{1 + \frac{\tau_s}{\tau_0}} \right)^{-1} \quad (\text{A.5})$$

which gives us the expression in (24). Comparative statics follow by inspection. \square

A.3 Corollary 1: Over-reaction and Environment-dependence

Since $\frac{\gamma^{REE}}{\tilde{\gamma}} > 1$, (46) shows that $\frac{\partial P^{PET}}{\partial s} > \frac{\partial P^{REE}}{\partial s}$, proving that PET leads to over-reaction to news relative to the rational benchmark. The wedge between PET and REE is then governed by $\frac{\gamma^{REE}}{\tilde{\gamma}}$:

$$\frac{\gamma^{REE}}{\tilde{\gamma}} = \frac{\tau_s}{\tau_s + \tau_0} \frac{\phi\tau_s + \tau_0}{\phi\tau_s} = \frac{\phi\tau_s + \tau_0}{\phi(\tau_s + \tau_0)} \quad (\text{A.6})$$

which is decreasing in ϕ and $\frac{\tau_s}{\tau_0}$.

For the individual level, notice that (22) shows that the sensitivity of the extracted signal with respect to the true one is given by:

$$\frac{\partial \tilde{s}}{\partial s} = 1 + \frac{\left(\frac{1-\phi}{\phi^2}\right) \tau_0}{\tau_s + \tau_0 - \left(\frac{1-\phi}{\phi}\right)^2 \tau_0} \quad (\text{A.7})$$

which is also decreasing in ϕ and $\frac{\tau_s}{\tau_0}$. □

A.4 Proposition 2: Amplification

The equilibrium signal PET agents extract from prices, \tilde{s}^{PET} , can be written as:³⁹

$$\tilde{s}^{PET} = s + \frac{\frac{\alpha}{\tilde{\gamma}} - \left(1 - \frac{\beta}{\tilde{\gamma}}\right)}{1 - \frac{\beta}{\tilde{\gamma}}} \left(s - (\mu_0 - AZ\tau_0^{-1})\right) \quad (\text{A.8})$$

By assumption, α and β are bounded below and above by strictly positive values.⁴⁰ It follows that $\tilde{\gamma}$ is bounded above too since $\tilde{\gamma} < \alpha + \beta$. Since s is fixed, these conditions ensure that $\lim_{\beta/\tilde{\gamma} \rightarrow 1} \tilde{s}^{PET} - s = \infty$. Turning to the wedge between PET and REE sensitivities of the equilibrium price with respect to the true signal s :⁴¹

$$\frac{\partial P^{PET}}{\partial s} - \frac{\partial P^{REE}}{\partial s} = \left(\frac{\alpha}{1 - \frac{\beta}{\tilde{\gamma}}}\right) - (\alpha + \beta) = \frac{\frac{\alpha}{\tilde{\gamma}} - \left(1 - \frac{\beta}{\tilde{\gamma}}\right)}{1 - \frac{\beta}{\tilde{\gamma}}} \quad (\text{A.9})$$

³⁹To see this, one can either rely on the derivation in Section 3, or simply compare the expression in (A.8) with (22), after having substituted in for the values of $\alpha = \frac{\phi\tau_s}{\tau_s + \tau_0}$, $\beta = \frac{(1-\phi)\tau_s}{\tau_s + \tau_0}$, and $\tilde{\gamma} = \frac{\phi\tau_s}{\phi\tau_s + \tau_0}$.

⁴⁰Formally, there exists $\underline{\phi} > 0$, $\underline{\tau}_s > 0$, $\bar{\tau}_s > 0$ and $\underline{\tau}_0 > 0$ such that $\underline{\phi} \leq \phi$, $\underline{\tau}_s \leq \tau_s \leq \bar{\tau}_s$, and $\underline{\tau}_0 \leq \tau_0$ when we vary $\beta/\tilde{\gamma}$. This is simply to ensure that $\alpha/\tilde{\gamma}$ does not go to 0.

⁴¹This wedge is also related to the wedge between $\tilde{s}^{PET} - s = \frac{1}{\beta} \left(\frac{\partial P^{PET}}{\partial s} - \frac{\partial P^{REE}}{\partial s}\right) (s - (\mu_0 - AZ\tau_0^{-1}))$.

so that this wedge is also *i*) increasing in $\beta/\tilde{\gamma}$, and *ii*) $\lim_{\beta/\tilde{\gamma} \rightarrow 1} \frac{\partial P^{PET}}{\partial s} - \frac{\partial P^{REE}}{\partial s} = \infty$. Next, we can express the wedge between PET and REE prices as:

$$P^{PET} - P^{REE} = \beta(\tilde{s}^{PET} - s) \quad (\text{A.10})$$

Since β is bounded below by a strictly positive number (by assumption), it follows immediately from the discussion above that this wedge verifies $\lim_{\beta/\tilde{\gamma} \rightarrow 1} P^{PET} - P = \infty$. Finally, the wedge in expected returns (defined in A.20),

$$\mathbb{E}_U[R]^{PET} - \mathbb{E}_U[R]^{REE} = \frac{\tau_s}{\tau_s + \tau_0}(\tilde{s} - s) - (P^{PET} - P^{REE}) \quad (\text{A.11})$$

$$= (\alpha + \beta)(\tilde{s} - s) - \beta(\tilde{s} - s) = \alpha(\tilde{s} - s) \quad (\text{A.12})$$

where the second equality uses the fact that $\frac{\tau_s}{\tau_s + \tau_0} = \alpha + \beta$ and the result in (A.10). Similarly,

$$\mathbb{E}_I[R]^{PET} - \mathbb{E}_I[R]^{REE} = -(P^{PET} - P^{REE}) = -\beta(\tilde{s} - s) \quad (\text{A.13})$$

where the second equality uses the result in (A.10). Therefore, (A.12) and (A.13) confirm that the limit property hold for the discrepancy between PET and REE expected and realized returns as well: $\lim_{\beta/\tilde{\gamma} \rightarrow 1} \mathbb{E}_i[R]^{PET} - \mathbb{E}_i[R]^{REE} = \infty$ for $i \in \{I, U\}$, since α and β are bounded. This also holds for realized returns as, on average, realized returns are equal to informed agents' expected returns. \square

A.5 Proof of Corollary 2: PET and Asset Pricing Puzzles

Over-reaction. The sensitivities of the equilibrium price to the true signal, across the two equilibrium concepts, can be rewritten as:

$$\frac{\partial P^{REE}}{\partial s} = \frac{\tau_s}{\tau_s + \tau_0}, \quad \frac{\partial P^{PET}}{\partial s} = \frac{\tau_s}{\tau_s + \tau_0 - \left(\frac{1-\phi}{\phi}\right)^2 \tau_0}. \quad (\text{A.14})$$

When the PET equilibrium is stable, the denominator of the PET sensitivity in (A.14) is positive (see Proposition 1). For $\phi \in (0, 1)$, it immediately follows that:

$$\frac{\tau_s}{\tau_s + \tau_0} < \frac{\tau_s}{\tau_s + \tau_0 - \left(\frac{1-\phi}{\phi}\right)^2 \tau_0}. \quad (\text{A.15})$$

Trading Volume. Trading volume for $K = \{REE, PET\}$ is defined as:

$$V^K = \phi|Z - X_I^K| + (1 - \phi)|Z - X_U^K| = 2\phi(1 - \phi)|X_U - X_I|, \quad (\text{A.16})$$

where the last equality uses the market clearing condition, $\phi X_I + (1 - \phi)X_U = Z$.

We make use of the fact that both informed and uninformed agents' demand functions use the same precision, so that (A.16) simplifies to:

$$V^K = \frac{2\phi(1 - \phi)\tau_s}{A} |\tilde{s}^K - s|. \quad (\text{A.17})$$

Trading volume is thus proportional to $|\tilde{s} - s|$, which is the amount of belief disagreement between I and U agents. The rational volume is 0. Substituting the PET extracted signal in (22) into equation (A.17) yields an expression of volume as a function of fundamentals:

$$V^{PET} = \frac{2(1 - \phi)^2 \tau_s \tau_0}{A\phi \left(\tau_s + \tau_0 - \left(\frac{1-\phi}{\phi}\right)^2 \tau_0 \right)} |s - \mu_0 + AZ\sigma_0^2| \quad (\text{A.18})$$

$$= \frac{2\phi(1 - \phi)\tau_s}{A} \left(\frac{\frac{\alpha}{\tilde{\gamma}} - \left(1 - \frac{\beta}{\tilde{\gamma}}\right)}{1 - \frac{\beta}{\tilde{\gamma}}} \right) |s - \mu_0 + AZ\sigma_0^2| \quad (\text{A.19})$$

First, notice that when $s = \mu_0 - AZ\sigma_0^2$, trading volume is zero (at this signal, the true model and the misspecified one are equal). Moreover, from (A.19), it is clear that V^{PET} is increasing in the size of the feedback effect $\frac{\beta}{\tilde{\gamma}}$, and that $\lim_{\frac{\beta}{\tilde{\gamma}} \rightarrow 1} V^{PET} = \infty$ when $\alpha/\tilde{\gamma}$ is not going to 0.

Returns. Expected returns are defined additively, for $i \in \{I, U\}$, and $K \in \{PET, REE\}$:

$$\mathbb{E}_i[R]^K = \mathbb{E}[v|\mathcal{I}_i^K] - P^K = \mathbb{E}[v|\mathcal{I}_i^K] - \left(\bar{\mathbb{E}}[v] - \bar{\mathbb{V}}[v]AZ\right) \quad (\text{A.20})$$

where $\bar{\mathbb{E}}[v]$ and $\bar{\mathbb{V}}[v]$ are weighted averages of agents' beliefs and precisions, respectively.⁴²

Since I agents have the same beliefs across REE and PET, differences in I agents' expected returns across equilibrium concepts are driven by differences in equilibrium prices. After some algebra, I agents' expected returns conditional on a signal s are given by:

$$\mathbb{E}_I[R]^{REE} = AZ(\tau_s + \tau_0)^{-1} \quad (\text{A.21})$$

$$\mathbb{E}_I[R]^{PET} = -\frac{(1-\phi)\tau_s}{\tau_s + \tau_0} \left(\frac{\left(\frac{1-\phi}{\phi^2}\right)\tau_0}{\tau_s + \tau_0 - \left(\frac{1-\phi}{\phi}\right)^2\tau_0} \right) \left(s - \left(\mu_0 - AZ\tau_0^{-1} \right) \right) + AZ(\tau_0 + \tau_s)^{-1} \quad (\text{A.22})$$

Since I agents' beliefs are rational, their expected returns are equal to average realized returns, conditional on s . The corresponding expected returns of U agents are given by:

$$\mathbb{E}_U[R]^{REE} = AZ(\tau_s + \tau_0)^{-1} \quad (\text{A.23})$$

$$\mathbb{E}_U[R]^{PET} = \frac{\phi\tau_s}{\tau_s + \tau_0} \left(\frac{\left(\frac{1-\phi}{\phi^2}\right)\tau_0}{\tau_s + \tau_0 - \left(\frac{1-\phi}{\phi}\right)^2\tau_0} \right) \left(s - \left(\mu_0 - AZ\tau_0^{-1} \right) \right) + AZ(\tau_0 + \tau_s)^{-1} \quad (\text{A.24})$$

The comparative statics in Proposition 2 follows by inspection. Since the denominator in (A.24) and (A.22) is positive when the equilibrium is stable, $\frac{\partial \mathbb{E}_U[R]^{PET}}{\partial s} > 0$ and $\frac{\partial R^{PET}}{\partial s} < 0$.

Inelastic Demands. This follows from Proposition 1. □

A.6 Proof of Proposition 3: Stability with General Misinference

The equilibrium is stable if and only if the aggregate demand for the risky asset is downward-sloping. The slope of the aggregate demand function can be expressed as:

$$-\frac{\partial X_{TOT}}{\partial P} = -\frac{1}{A} \left(1 - \frac{\hat{\beta}}{\hat{\gamma}} \right) P + constant \quad (\text{A.25})$$

leading to a downward sloping demand function and a stable equilibrium if and only if $\frac{\hat{\beta}}{\hat{\gamma}} < 1$.

⁴²Let $\tau_i = \mathbb{V}[v|I_i]$ for $i \in \{I, U\}$. Then $\bar{\mathbb{E}}[v] = \frac{\phi\tau_I}{\phi\tau_I + (1-\phi)g} \mathbb{E}[v|I] + \frac{(1-\phi)g}{\phi\tau_I + (1-\phi)g} \mathbb{E}[v|U]$, and $\bar{Var} = (\phi\tau_I + (1-\phi)g)^{-1}$. The fact that $P^K = \bar{\mathbb{E}}[v] - \bar{\mathbb{V}}[v]AZ$ then follows by inspection of (19) and (21).

A.7 Proof of Proposition 4: Correct Inference from Prices

This follows from (47), where $s = \tilde{s}$ if and only if $\hat{\alpha} + \hat{\beta} - \hat{\gamma} = 0$. When $\hat{\alpha} + \hat{\beta} - \hat{\gamma} > 0$ traders infer a signal which is more extreme than the true one, and conversely when $\hat{\alpha} + \hat{\beta} - \hat{\gamma} < 0$. \square

A.8 Proof of Proposition 5: Over/Under-reaction

This follows from inspecting (46) and comparing it to the rational benchmark:

$$P^{REE} = (\alpha + \beta)s + (1 - \alpha - \beta)\mu_0 \quad (\text{A.26})$$

The sensitivity of equilibrium prices is then higher than in the rational case when:

$$\frac{\hat{\alpha}}{1 - \frac{\hat{\beta}}{\tilde{\gamma}}} > \alpha + \beta \quad (\text{A.27})$$

which can be rearranged to yield Proposition 5. The under-reaction case is symmetric. \square

A.9 Proof of Corollary 3: Mistakenly Assuming Rationality

In equilibrium, (46) shows that the sensitivity of prices to new information is given by:

$$\frac{\partial P}{\partial s} = \frac{\hat{\alpha}}{1 - \frac{\hat{\beta}}{\tilde{\gamma}}} \quad (\text{A.28})$$

Let $\hat{\gamma} \equiv \hat{\alpha} + \hat{\beta}$ and $\tilde{\gamma}^{REE} \equiv \alpha + \beta$. If $\hat{\alpha} \leq \alpha$ and $\hat{\beta} \leq \beta$ such that $\hat{\alpha} + \hat{\beta} < \alpha + \beta$. When this is the case, mislearning from fundamentals pushes towards under-reaction. Moreover, $\hat{\gamma} < \tilde{\gamma}^{REE}$, and from (A.28) we see that the sensitivity of prices to new information is lower when $\tilde{\gamma} = \tilde{\gamma}^{REE}$ than when $\tilde{\gamma} = \hat{\gamma}$. In other words, in this case mislearning from prices contributes to even more under-reaction, thus amplifying the original bias in mislearning from fundamentals. The converse is true when $\hat{\alpha} \geq \alpha$ and $\hat{\beta} \geq \beta$ such that $\hat{\alpha} + \hat{\beta} > \alpha + \beta$. Finally, the equilibrium is unstable when $\tilde{\gamma} < \hat{\beta}$, which cannot happen when $\tilde{\gamma} = \hat{\gamma}$ since $\hat{\alpha} > 0$, but can happen when $\tilde{\gamma} = \tilde{\gamma}^{REE}$ if $\alpha + \beta < \hat{\beta}$. Therefore, when $\hat{\alpha} + \hat{\beta} > \alpha + \beta$, misinference from mistakenly assuming the world is rational can make the equilibrium unstable. \square

B Derivation of the REE Inference Problem

To construct the mapping that uninformed agents use to extract information from prices, we need to determine the market clearing condition which they think is generating the price that they observe: $\phi\tilde{X}_I + (1 - \phi)\tilde{X}_U = Z$. This, in turn, requires us to specify what uninformed traders think about other agents' posterior beliefs. Specifically, they think informed traders trade on the signal \tilde{s} :

$$\mathbb{E}[v|\tilde{\mathcal{I}}_U] = \frac{\tau_s}{\tau_s + \tau_0}\tilde{s} + \frac{\tau_0}{\tau_s + \tau_0}\mu_0 \quad (\text{B.29})$$

$$\mathbb{V}[v|\tilde{\mathcal{I}}_U] = (\tau_s + \tau_0)^{-1} \quad (\text{B.30})$$

Moreover, under rational expectations uninformed traders realize that other uninformed traders are also learning from prices. Let us guess that they think other uninformed traders use the following linear mapping to infer information from prices: $P = \gamma_{-i}\tilde{s} + (1 - \gamma_{-i})(\mu_0 - AZ\tau_0^{-1})$. This translates into the following posterior beliefs:

$$\mathbb{E}[v|\tilde{\mathcal{I}}_U] = \frac{\tau_s}{\tau_s + \tau_0} \left(\frac{1}{\gamma_{-i}}P - \frac{1 - \gamma_{-i}}{\gamma_{-i}}(\mu_0 - AZ\tau_0^{-1}) \right) + \frac{\tau_0}{\tau_s + \tau_0}\mu_0 \quad (\text{B.31})$$

$$\mathbb{V}[v|\tilde{\mathcal{I}}_U] = (\tau_s + \tau_0)^{-1} \quad (\text{B.32})$$

Market clearing then leads to the following perceived price function:

$$P = \left(\frac{\frac{\tau_s}{\tau_s + \tau_0}}{1 - (1 - \phi) \left(\left(\frac{\tau_s}{\tau_s + \tau_0} \right) \frac{1}{\gamma_{-i}} - 1 \right)} \right) \tilde{s} + \left(1 - \frac{\frac{\tau_s}{\tau_s + \tau_0}}{1 - (1 - \phi) \left(\left(\frac{\tau_s}{\tau_s + \tau_0} \right) \frac{1}{\gamma_{-i}} - 1 \right)} \right) (\mu_0 - AZ\tau_0^{-1}) \quad (\text{B.33})$$

Given their beliefs about other agents, uninformed traders would then invert this mapping to infer information from prices. However, common knowledge of rationality implies that all uninformed traders reason in this way, and that $\gamma_i = \gamma_{-i} = \gamma_{REE}$, leading to the following fixed-point:

$$\gamma_{REE} = \frac{\frac{\tau_s}{\tau_s + \tau_0}}{1 - (1 - \phi) \left(\left(\frac{\tau_s}{\tau_s + \tau_0} \right) \frac{1}{\gamma_{REE}} - 1 \right)} \implies \gamma_{REE} = \frac{\tau_s}{\tau_s + \tau_0} \quad (\text{B.34})$$

Therefore, under common knowledge of rationality uninformed traders infer information from prices by inverting the following mapping, which coincides with the mapping we derived in (14) in the text:

$$P_{REE}^{Mis}(\tilde{s}) = \frac{\tau_s}{\tau_s + \tau_0} \tilde{s} + \frac{\tau_0}{\tau_s + \tau_0} (\mu_0 - AZ\tau_0^{-1}) \quad (\text{B.35})$$