Learning, Over-reaction and the Wisdom of the Crowd*

Daniele d'Arienzo[a] and Matteo Bizzarri[b]

[a]Università Bocconi, Department of Finance

[b]Università Bocconi, Department of Economics, and CSEF

August 20, 2021

Abstract

How do departures of individual beliefs from Bayesian rationality impact welfare and informational efficiency in financial markets? We allow agents posterior beliefs to depart from the rational expectations hypothesis and to over/under-react to news (i.e. to changes in fundamentals and/or prices). We find that both in a financial market with sequential trading and unit demand and supply, as well as in a financial market with simultaneous trading, with volume and unbounded signals, over-reaction has a positive informational externality, while under-reaction exacerbates any informational friction. Over-reaction mitigates individual losses due non-Bayesian beliefs. Because agents over-react to news, prices aggregate more efficiently private information, relative to the rational case. As a consequence, over-reaction decreases the likelihood of informational cascades, increases the informational content of prices, and it increases welfare.

*We wish to thank for useful comments Nicola Gennaioli, Philip Matejka, Filippo Massari, Alex Teytelboym, Fernando Vega-Redondo, and seminar participants in Bocconi.

1 Introduction

Traders in financial markets constantly update their beliefs about valuations of financial assets, as a consequence of changes in market prices, fundamentals and/or investment choices of other traders. Bayesian rationality prescribes a certain law of motion for the dynamics of such beliefs, which, however, has been found at odds with reality (Benjamin (2019)). Individual departures from Bayesian rationality in the forms of under/over-reaction to information have been documented empirically (Bordalo et al. (2018)). Moreover, such departures explains several facts about macro-financial variables, such as credit cycles (Bordalo et al. (2018)), stock returns (Bouchaud et al. (2019), Bordalo et al. (2018)), interest rates (d'Arienzo (2020)) and even the likelihood of a financial crisis (Maxted (2019)). How do departures of individual beliefs from Bayesian rationality impact welfare and informational efficiency in financial markets?

To answer this question, we follow the diagnostic expectations model (Bordalo et al. (2016) and Bordalo et al. (2018)), and we define over/under-reaction in beliefs as a one-parameter deviation from Bayesian updating. The model is parsimonious, tractable and beliefs are forward looking. Biased beliefs depart from Bayesian rationality by under/over-reacting to recent information. Our model applies the same logic of the diagnostic expectations model of (Bordalo et al. (2018)) to a learning setting.[1] Diagnostic learning converges to true fundamentals, as in the Bayesian case, but beliefs over/under-react to recent data.

Then, we embed diagnostic learning into two classic workhorse models of trade under asymmetric information: the sequential trading model (Glosten and Milgrom (1985)), and the simultaneous linear quadratic Gaussian trading model (used e.g. in Grossman and Stiglitz (1980)).

In the former setting, agents enter sequentially the market, observe past trades of other agents, receive a binary informative signal about fundamentals, take a buy or sell action (for the single unit of the traded asset) and finally exit the market. A competitive market-maker efficiently sets prices as expected values of fundamentals given public information. Under asymmetric information between the market maker and the traders, such a model features

[1] A different application of diagnostic learning in financial markets is investigated by (Santosh (2020)).

an informational inefficiency. When traders are ex-ante much more informed about asset valuations relative to the market maker, then all agents agents buy, regardless of receiving a good or bad private signal about the fundamentals. As a consequence, the market maker cannot discriminate private information from trades, and private information is not efficiently incorporated into prices. We show that over-reaction mitigates the phenomenon. When agents over-react, those who received are overly pessimistic, their valuation is below the valuation of the market maker, hence they sell. Then, the market maker infers their private bad signal from this trade, and next period's market price efficiently adjusts accordingly. In the long run, the price get closer to the fundamental, relative to a market populated by Bayesian agents.

We also allow also the market maker to depart from rationality and to under/over-react. In this case, the market informational efficiency depends on the relative mis-reaction between traders and the market maker, and it is time dependent.

The limit case where the market maker has the maximum level of under-reaction possible is interesting in itself, because in this setting prices are fixed over time, and so, from the perspective of future agents, buy and sell orders have a purely informational content. Hence, we recover a pure social learning setting, analogous to (Bikhchandani et al. (1992)) and (Banerjee (1992)). We show in this case there is a finite level of over-reaction that maximizes the probability that the society learns the fundamental, that in this setting can roughly be translated as "there is a cascade on *buy* if the asset is good, a cascade on *sell* if the asset is bad".

Crucially, the informational role of over-reaction does not rely on the presence of cascades. Indeed, we show that the same force applies in a simultaneous trading model, with mean-variance preferences (Glosten and Milgrom (1985)). In such a model, in equilibrium, the price can be seen as a signal of the fundamental and the precision of such signal is *increasing* in the degree of over-reaction, as first shown in (Bordalo et al. (2020)). We show that with over-reaction, agents trade more aggressively for the same signal, hence they increase the informativeness of the price, which, in turn, makes agents better off. At the same time, larger trades penalizes individual utilities of traders, due the risk aversion. We find the first effect markedly dominates. Social welfare is increasing in the level of over-reaction, whose socially optimal value is therefore arbitrarily large: in this scenario, beliefs of traders concentrate of their observed private and price signals. Agents aggressively trade based on those signals, which, in turn, generates high price informativeness.

Literature Review

This work contributes to two streams of literature: the literature on herding and cascades in financial markets, and the literature on under/over-reaction of beliefs in financial markets.

The literature on herding and informational cascades (Banerjee (1992) and Bikhchandani et al. (1992)) in financial markets study studies weather market prices can ameliorate such informational frictions. (Avery and Zemsky (1998)) shows that in a binary state space herding and informational cascades occur where there sufficiently many sources of uncertainty; (Park and Sabourian (2011)) shows that cascades can happen if the state space is more complex. We depart from these papers in two dimensions: first, we explore the consequences of non-Bayesian agents that display under/over-reaction. Second, in a model with no cascades, to generate cascades, we take a simpler road than the aforementioned papers: we depart from the common prior assumption by assuming sufficiently high ex-ante asymmetric information.

The literature on diagnostic expectations explains various macro-financial regularities: Bordalo et al. (2018), Bordalo et al. (2018), Bordalo et al. (2019), Bordalo et al. (2020). While these papers matches empirical regularities, we investigate the welfare properties of the same departure from Bayesian rationality. Walther (2020) explores macro-prudential policy implications with extrapolative beliefs.

Our work is also connected to the literature on non Bayesian learning: examples that study social learning problems are Bohren and Hauser (2021), Molavi et al. (2018), Frick et al. (2020). All these papers focus on long run convergence properties of beliefs, not informational efficiency of markets. The closest to ours is Santosh (2020) that studies Kalman filtering and stochastic volatility under learning analog of diagnostic expectations similar to ours. Finally, our work is connected to the literature exploring the informational efficiency of markets, a large body of literature whose full review is beyond our scopes. Here we limit to note the recent work by Vives (2017), exploring the inefficiencies in information transmission in a linear quadratic setting similar to the model we explore in Section 4.

The paper is organized as follows: Section 2 lays out the definition of the diagnostic updating we use throughout the paper as long as a few basic properties. Section 3 presents the analysis of a sequential trading model under diagnostic learning, as well as the particular case in which the market maker features extreme under-reaction and the model becomes analogous to a pure social learning setting. Section 4 makes the point that the possibility of cascades is not the driving force of our results, by showing that over-reaction is still welfare improving in a classic linear quadratic model of trade with unbounded signals. All proofs are in the Appendix.

2 A non Bayesian learning model

In this section we introduce a learning model based on diagnostic expectations. Consider an agent that has to *learn* the state of the world ω, on which she has a prior belief, with density $p_0(\omega)$. The agent observes a signal X, whose likelihood is known to be $l(X|\omega)$. The standard Bayesian updating operator takes as inputs the prior density, the likelihood function and the observed signal, and it prescribes to update beliefs about ω from $p_0(\omega)$ to the Bayesian posterior:

$$\mathcal{BU}(l, p_0)(\omega) := \frac{l(X|\omega)p_0(\omega)}{\int l(X|\omega')p_0(\omega')d\omega'}. \tag{1}$$

Following Bordalo et al. (2018), we assume that agent's updating rule departs from the standard Bayes rule, as follows:

$$\mathcal{BU}^\theta(l, p_0)(\omega) = \frac{1}{Z}\mathcal{BU}(l, p_0)(\omega)\left(\frac{\mathcal{BU}(l, p_0)(\omega)}{p_0(\omega)}\right)^\theta, \tag{2}$$

where Z is a normalization constant. This rule can be founded on the *representativeness heuristics* (Tversky and Kahneman (1974)): current information is exaggerated when for representative states ω, which are those whose the likelihood ratio relative past assessment, $\frac{\mathcal{BU}(l,p_0)(\omega)}{p_0(\omega)}$, increased. The scalar parameter $\theta > -1$ controls the departure from the Bayesian case. When $\theta > 0$ the model delivers *over-reaction* to information. Specifically, the previous formula says that states ω which are more likely under $\mathcal{BU}(l, p_0)(\omega)$ than under $p_0(\omega)$, i.e. *representative* states, are over-weighted. On the contrary, states ω which are less likely under $\mathcal{BU}(l, p_0)(\omega)$ than under $p_0(\omega)$, are under-weighted. Thus, we say that for $\theta > 0$ posterior beliefs over-react to information. On the contrary, for $-1 < \theta < 0$, posterior beliefs under-react to information.

The next proposition sums up two general properties of diagnostic updating that are useful to keep in mind in the following. First, *consistency*: when facing multiple observations, $X_1, \ldots, X_t$, agents can update their beliefs in one round, or sequentially; in this latter case, at each step, the prior belief is the previous step distorted posterior belief:

$$\mathcal{BU}_t^\theta(l, p_0)(\omega) = \frac{l(X_1, \ldots X_t|\omega)^{1+\theta} p_0(\omega)}{\int l(X_1, \ldots X_t|\omega')^{1+\theta} p_0(\omega') d\omega'}.$$

Consistency states that the posterior distribution after observing many signals is the same if we perform the updating once simultaneously, or sequentially (hence, also independent of the order). This is a convenient property when we study the dynamics. Second: diagnostic updating does not affect long-run learning by an agent in isolation.

Theorem 1. *Diagnostic updating satisfies the following properties:*

1. Consistency *For $k \in \{1, \ldots, t-1\}$:*

$$\mathcal{BU}_t^\theta(l, p_0)(\omega) = \mathcal{BU}_{t-k}^\theta(l, \mathcal{BU}_k^\theta(l, p_0))(\omega)$$

2. Asymptotic learning (in isolation) *Consider i.i.d. signals X_t, with likelihood $l(X_t|\omega)$, such that $l(X|\omega) \neq l(X|\omega')$ if $\omega \neq \omega'$, and consider a prior p_0 on ω. Consider the cases where the prior is discrete or both signals and prior are continuous (for now).*

   Call l_ω^∞ the measure on the space of realizations $X = (X_1, X_2, \ldots)$ defined by the product measure of the likelihoods $l(\cdot|\omega)$. Then for p_0-almost every ω the diagnostic learning converges a.s., that is:

$$P^\theta(\omega \in B|X_1, \ldots, X_t) \to 1_{\omega^* \in B} \quad l_{\omega^*}^\infty - a.s.$$

3 Sequential Learning and Market Efficiency

How does over-reaction impact informational efficiency in financial markets? Even in rational financial markets, informational efficiency may be constrained by herding, namely the phenomenon where traders temporarily neglect private information and copy actions of other traders. (Lakonishok et al. (1992)) and (Nofsinger and Sias (1999)) document herding using trading data of institutional investors; Cipriani and Guarino (2014) find herding an experimental setting with sequential trading. Theoretically, (Avery and Zemsky (1998)) builds on the financial market model of (Glosten and Milgrom (1985)) and show that rational herding occurs when there are multiple dimensions of uncertainty about the true fundamental value of the single asset traded in the market. (Park and Sabourian (2011)) consider rich structures of signals and state spaces, and characterize the occurrence of rational herding in financial markets with sequential trading depending on the shape of the signals. Both the papers also characterize the occurrence of informational cascades, which are defined as permanent herding. Informational cascades have broader implications than financial markets: they can occur also in the purely sequential social learning settings, such as (Banerjee (1992)) and (Bikhchandani et al. (1992)).

Here, we consider a simple model of financial market with sequential trading and we show that over-reaction may increase market informational efficiency, as opposed to Bayesian rationality. Because traders over-react to their signals, we show that with over-reaction they inject more information into prices, and it is therefore less likely to get stuck into informational cascades, relative to Bayesian rationality. Traders' valuation of the asset is more extreme with over-reaction, hence buy trade is more likely to signal positive information about the asset. On the contrary, under-reaction shrink traders' valuation toward the public valuation, and with under-reaction it is less likely to escape an informational cascade, relative to Bayesian rationality. This is also true in the case of purely social learning, which is the limit of our model when prices are infinitely sticky. When traders react differently to information relative to the market maker, the model has richer predictions: over-reaction exaggerate private signals, yet differential reactions across traders and the marker makers also distorts the weight attached to public information. The balance between the two forces depends on the details of the market and determines the actual features of the price process.

3.1 The Model

A single asset with binary fundamental value $V \in \{0, 1\}$ is traded. Time is discrete. At each time $t \geq 1$, a risk-neutral market maker (e.g. a broker or the stock exchange) sets the price of single traded unit of the asset as the expected value of the fundamental, given public information and given his beliefs. Public information is defined as the history of past trades: $H_t = (a_1, \ldots, a_t)$, where $a_t = B$ (Buy) or $a_t = S$ (Sell). Therefore, the price of the asset at time t reads:

$$p_t = \mathbb{E}_{MM}[V|H_t] = Pr_{MM}(V = 1|H_t), \tag{3}$$

where the label "MM" denotes the expectation is taken using the beliefs of the market maker. At each time $t \geq 1$, a risk-neutral trader enters the market, she trades and the she exits the market forever. She observes both public information, H_t, and an informative binary private signal $s_t \in \{-1, 1\}$ about the fundamental. She buys if and only if:

$$\mathbb{E}_{TR}[V|H_t, s_t] \geq p_t, \tag{4}$$

otherwise she sells. "TR" denotes the expectation is taken using the beliefs of the trader. The model features asymmetric information: traders observe private signals, while the market maker does not. Indeed the goal the seminal paper of (Glosten and Milgrom (1985)), which started the sequential trading literature, is to study the effects of asymmetric information in financial markets. We assume further that such asymmetry in information occurs also time $t = 0$: the market maker has no information about the fundamental value, namely $Pr_{MM}(V = 1) = \frac{1}{2}$; traders, on the contrary, have some prior information, $Pr_{TR}(V = 1) = p > \frac{1}{2}$ (with no loss of generality about the direction of the inequality). As we shall show, such ex-ante asymmetric information generates the possibility of informational cascades in the rational version of the model, while at the same time this specification keeps the model extremely tractable. Different ways of generating herding and cascades include: additional sources of uncertainty about the fundamental value of the asset (Avery and Zemsky (1998)), richer action-state space structures (Park and Sabourian (2011)), and risk-aversion (Decamps and Lovo (2006)).

A rational and risk-neutral trader at time t, on top of the public information H_t, has private information: she receives a signal s_t, which is informative about the fundamental value of the asset: $Pr(s_t = \pm 1 | 2V - 1 = \pm 1) = q$ and $Pr(s_t = \pm 1 | 2V - 1 = \mp 1) = 1 - q$, with $q > \frac{1}{2}$. Consider the valuation problem of the trader. If she receives bad news, $s_t = -1$, then her valuation of the asset reads:

$$\mathbb{E}_{TR}\left[V | H_t, s_t = -1\right] = Pr_{TR}(V = 1 | H_t, s_t = -1) \tag{5}$$

$$= \frac{Pr_{TR}(V = 1 | H_t) Pr(s_t = -1 | V = 1, H_t)}{Pr_{TR}(s_t = -1 | H_t)} = \frac{1}{\pi_t + (1 - \pi_t)\frac{q}{1-q}} \pi_t, \tag{6}$$

where $1 - q = Pr(s_t = -1 | V = 1, H_t) = Pr(s_t = -1 | V = 1)$ because signals are conditionally i.i.d., and $\pi_t := \mathbb{E}_{TR}[V | H_t]$. Similarly, If she receives good news, $s_t = 1$, then her valuation of the asset is:

$$\mathbb{E}[V | H_t, s_t = 1] = \frac{1}{\pi_t + (1 - \pi_t)\frac{1-q}{q}} \pi_t. \tag{7}$$

On the contrary the market maker evaluates the asset as:

$$\mathbb{E}[V | H_t] = p_t = \frac{1}{1 + \frac{Pr(H_t | V=0)}{Pr(H_t | V=1)}}. \tag{8}$$

$Pr(H_t | V = 1)$ and $Pr(H_t | V = 0)$ have no labels "MM" because the likelihood of a stream of trades H_t is independent of the prior of the market maker. Following (Avery and Zemsky (1998)), we say there is no herding at time t if $\mathbb{E}_{TR}[V | H_t, s_t = 1] > p_t > \mathbb{E}_{TR}[V | H_t, s_t = -1]$. Without herding, a trader who receives a good signal buys, a trader who receives a bad signal sells: observed trades are then informative about private signals, and the learning process of the market maker incorporate them into prices. Prices efficiently adjust in the rational case to prevent herding.

Does herding occur under diagnostic expectations? Herding does not occur if and only if:

$$\mathbb{E}[V|H_t, s_t] > p_t > \mathbb{E}[V|H_t, s = -1] \tag{9}$$

$$\iff \frac{Pr_{TR}(V|H_t, s_t = 1)}{1 - Pr_{TR}(V|H_t, s_t = 1)} > \frac{p_t}{1 - p_t} > \frac{Pr_{TR}(V|H_t, s_t = -1)}{1 - Pr_{TR}(V|H_t, s_t = -1)} \tag{10}$$

$$\iff \frac{Pr(H_t|V = 1)}{Pr(H_t|V = 0)} \frac{p}{1 - p} \frac{q}{1 - q} > \frac{Pr(H_t|V = 1)}{Pr(H_t|V = 0)} > \frac{Pr(H_t|V = 1)}{Pr(H_t|V = 0)} \frac{p}{1 - p} \frac{1 - q}{q} \tag{11}$$

$$\iff \frac{q}{1 - q} > \frac{1 - p}{p} > \frac{1 - q}{q}. \tag{12}$$

If prior beliefs of traders are not too extreme, then for all trading histories H_t, herding does not occur and prices are perfectly informative about private information. On the contrary, there is a buy(sell)-cascade starting from time $t = 1$ if p is close enough to zero (one): then, prices are completely uninformative about private information.

While in this setting price informativeness can either be full or completely absent, our results do not depend on the sequential nature of the trading, as discussed in Section 4.

How does diagnostic learning impact price informativeness? We call θ the distortion parameter of traders and $\bar{\theta}$ the distortion parameter of the market maker. Consider first the homogeneous case $\theta = \bar{\theta}$. There is no herding if and only if:

$$\mathbb{E}^{\theta}_{TR}[V|H_t, s_t = 1] > p_{\theta,t} := \mathbb{E}^{\theta}_{MM}[V|H_t] > \mathbb{E}^{\theta}_{TR}[V|H_t, s_t = -1] \tag{13}$$

$$\iff \left(\frac{q}{1 - q}\right)^{1+\theta} > \frac{1 - p}{p} > \left(\frac{1 - q}{q}\right)^{1+\theta}. \tag{14}$$

over-reaction ($\theta > 0$) prevents herding, relative to the Bayesian-rational case ($\theta = 0$), in the sense that the condition to avoid herding is less stringent. On the contrary, with under-reaction ($\theta < 0$) the condition to avoid herding is more stringent.

Second, consider the case where θ is the distortion of the traders, while $\bar{\theta}$ is the distortion of the market maker. The following results characterizes the occurrence of herding.

Proposition 3.1. *There is no herding in the model at time t, if and only if over-reaction of traders to the private signal s_t dominates on the relative mis-reaction of traders with respect* *to the market maker to public information H_t:*

$$\left(\frac{q}{1-q}\right)^{1+\theta}\left(\frac{Pr(H_t|V=1)}{Pr(H_t|V=0)}\right)^{\theta-\bar{\theta}} > \frac{1-p}{p} > \left(\frac{1-q}{q}\right)^{1+\theta}\left(\frac{Pr(H_t|V=1)}{Pr(H_t|V=0)}\right)^{\theta-\bar{\theta}}. \tag{15}$$

The proof simply follows from traders' and market maker valutation formulas. Distorted learning has three effects. First, traders distortions θ over-weight the private signal s_t, exactly as in the case $\theta = \bar{\theta}$. Second, learning from public information H_t is distorted because the frequency of actions H_t differs relative to the rational case: with over-reaction buying is more likely after good news. Third, learning from public information is distorted. If traders have sufficiently high θ relative to the market maker, namely if $\theta - \bar{\theta} > 0$, then traders over-react also to public information, relative to the market maker. If $\theta - \bar{\theta} < 0$ instead, traders under-react to public information. The latter condition can happen even in a world where both traders and the market maker over-reacts, because what matters for distorted learning from public information is the difference between the distortion of traders and that of the market maker. If signals are sufficiently informative, or if θ is sufficiently high, then over-reaction prevents herding. The dynamics of the system, in terms of price dynamics is rich, and a full characterization of it is beyond the scope of this paper. Our focus is instead on informational efficiency, and we shall consider two benchmark cases: a rational market-maker ($\bar{\theta} = 0$) and a perfectly rigid market-maker ($\bar{\theta} = -1$), where the latter maps our model into the purely social learning setting of (Banerjee (1992)) and (Bikhchandani et al. (1992)).

3.2 Rational Market Maker

The dynamics of the model is characterized as follows.

Theorem 2. *Assume $\bar{\theta} = 0$ and $\theta > \theta^* := \frac{\frac{p}{1-p}}{\frac{q}{1-q}} + 1$. Then: after $s_1 = 1$, there is a buy cascade; after $s_1 = s_2 = -1$, there is a sell cascade; finally, after $s_1 = -1, s_2 = 1$ the trader's choice at time $t = 3$ (given s_3), is identical to the trader's choice at time $t = 1$ (given $s_1 = s_3$).*

Moreover, the probability of a buy and sell cascades are:

$$Pr(Buy\ cascade|V=1) = Pr(Sell\ cascade|V=0) = \frac{q}{1-(1-q)q} \tag{16}$$

$$Pr(Sell\ cascade|V=1) = Pr(Buy\ cascade|V=0) = 1-\frac{q}{1-(1-q)q} \tag{17}$$

We measure market informational inefficiency from trader's perspective as the expected square distance between the actual price in the long-run, p_∞, and the fundamental value V:

$$\mathbb{E}[(V-p_\infty)^2]. \tag{18}$$

Such loss equals $p_\infty = \frac{1}{4}$ in the case of rational traders, namely $\theta = 0$. On the contrary, under the hypothesis of Theorem (2), the price under a buy cascade (given $V = 1$) is q, while the price under a sell cascade (given $V = 0$) is $(1-q)^2$ (see the proof of theorem (2)). The case $V = 0$ is symmetric. Hence the informational inefficiency reads:

$$\mathbb{E}[(V-p_\infty)^2] = (1-q)^2(Pr(\text{Buy cascade}|V=1) + (1-q)^4(Pr(\text{Sell cascade}|V=1)) < \frac{1}{4}. \tag{19}$$

Market informational efficiency improves with over-reaction ($\theta \geq \theta^*$). Over-reaction to news has the positive effect of partially healing negative informational externalities. While in this stylized setting over-reaction has only a positive effect, in general sequential trading models it may also induce herding (temporary cascades), as derived in proposition (3.1). The optimal amount of over-reaction in a more realistic market is therefore pinned by its specific features. Yet, the positive force of over-reaction on the diffusion of the information in financial markets, which is the focus of the current analysis, is a general effect. We now move to the social learning case.

3.3 Sequential Learning and Efficiency: the case of pure social learning

In this section, we apply our behavioral model of learning to the special case where $th\bar{e}ta =$ -1: the is no market, and the price of the asset is fixed and equal to $\frac{1}{2}$. This recovers the models of social learning models (Banerjee (1992)) and (Bikhchandani et al. (1992)). We now analyze the dynamics of the model.

Exactly as in the previous section, actions are perfectly revealing at time $t = 1$. To understand what are the implications of the distortion θ for cascades and learning, let us start with the following definition.

Definition 3.1. The *Informational efficient region* (IE) is the set of parameters given by the union of:

$$\theta + 1 \geq \frac{\ln \frac{p}{1-p}}{\ln \frac{1-q}{q}} \quad p \geq \frac{1}{2}, \tag{20}$$

and

$$\theta + 1 \leq \frac{\ln \frac{p}{1-p}}{\ln \frac{q}{1-q}} \quad p < \frac{1}{2}. \tag{21}$$

The Informational efficient region is the region of parameters such that time $t = 1$ traders buy if and only if $s_1 = 1$ and therefore "communicate" her private signal to future agents. Outside the efficient region, the trader instead chooses the action consistent with his prior regardless of her the signal. This is crucial in characterizing the behavior of the model. The following proposition describes such behavior.

Proposition 3.2. *If the parameters are in the Informational efficient region, then:*

If $p \geq \frac{1}{2}$ *If the first signal is 1, there is a cascade on 1. If the first two signals are* $(0, 0)$ *there is a cascade on 0. If the first two signals are* $(0, 1)$*, then the third agent faces the same problem of agent 1. The probability of learning is* $Pr(a_\infty = V) = \frac{pq+(1-p)q^2}{1-q(1-q)}$.

If $p < \frac{1}{2}$ *There is a cascade on 1 if the first signals are* $(1, 1)$*, there is a cascade on 0 if the* *first signal is 0. If the first two signals are* $(1,0)$*, then the third agent faces the same problem of agent 1. The probability of learning is* $Pr(a_\infty = V) = \frac{pq^2+(1-p)q}{1-q(1-q)}$.

If the parameters are outside of the Informational efficient region, then:

1. *If* $\frac{1}{2} > p$*, then all agents play 0 with probability 1 and the probability of learning is* $1-p$.

2. *If* $\frac{1}{2} \leq p$*, then then all agents play 1 with probability 1, and the probability of learning is* p.

By the form of the results, we can already see that a larger θ creates more room for learning, by enlarging the Informational efficient region. In the following, we make this argument formal.

A way to quantify the size of the parameter space is to think of the parameters p and q as drawn before the process starts. from a distribution μ, with full support on $(0,1) \times (\frac{1}{2}, 1)$. Denote $a_\infty = \lim_{t\to\infty} a_t$. Consider regions as $R_1 = IE \cup \{p > \frac{1}{2}\}$, $R_2 = IE \cup \{p \leq \frac{1}{2}\}$, $N_1 = \overline{IE} \cup \{p >$
$frac12\}$, and $N_2 = \overline{IE} \cup \{p \leq \frac{1}{2}\}$. Then the ex-ante probability of learning the correct state of the world is:

$$Pr(a_\infty = V) =$$
$$= \int \left(pI_{N_1} + (1-p)I_{N_2} + \frac{pq + (1-p)q^2}{1-q(1-q)} I_{R_1} + \frac{pq^2 + (1-p)q}{1-q(1-q)} I_{R_2} \right) \mathrm{d}\mu,$$

where $I(\cdot)$ represents the indicator function.

Let us define a level of θ *ex-ante efficient* if it achieves the maximum of this probability. The following figures illustrate the situation. In figure 1 we draw the region where Mr1 playing $a_1 = s_1$ is *socially* efficient, and the region where (Bayesian) Mr1 playing $a_1 = s_1$ is *individually* efficient (i.e. optimal). As is clear from the figure, a Bayesian updating rule does not maximize the learning probability: there is a region where it would be socially efficient that Mr1 plays $a_1 = s_1$, but a Bayesian agent, since he does not internalize the information externality on other agents, does not. In figure 2, we plot instead the informational efficient regions for different values of the parameter θ: we can see that there are moderate values

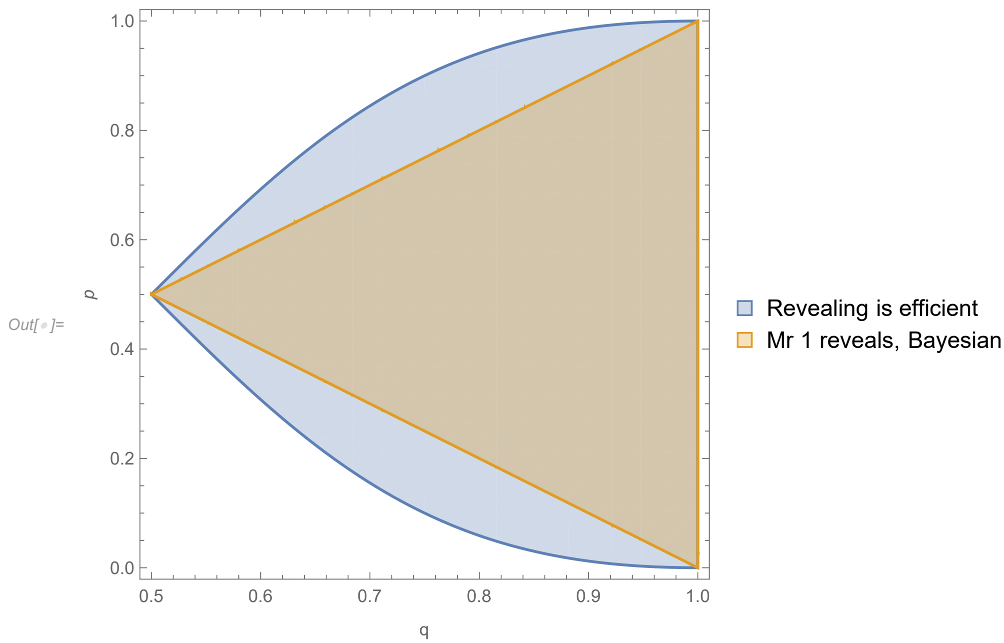

Figure 1: The areas of the parameter space where revealing is efficient, and the area where a Bayesian agent reveals. The figure shows that the Bayesian updating fails to be socially efficient: there is a region where agent 1 does not reveal but it would be socially optimal to do so.

of over-reaction that increases the probability of learning. In the following proposition, we show that there is a value of θ that actually achieve ex-ante efficiency.

Theorem 3. *If the parameters p, q are drawn from a distribution μ with full support on $(0,1) \times (\frac{1}{2}, 1)$, the distorted updating with $\theta = 1$ is ex-ante efficient.*

Given the importance of the result, we report the proof here in the main text.

Proof. Let us focus on the subset of the parameter space where $\{p > \frac{1}{2}\}$, the reasoning for $p < \frac{1}{2}$ is analogous. The the ex-ante probability of learning is:

$$\int \left(pI\left(\theta + 1 \leq \frac{\ln\left(\frac{p}{1-p}\right)}{\ln \frac{q}{1-q}}\right) + \frac{pq + (1-p)q^2}{1 - q(1-q)} I\left(\theta + 1 > \frac{\ln\left(\frac{p}{1-p}\right)}{\ln \frac{q}{1-q}}\right)\right) \mathrm{d}\mu.$$

In the Informational efficient region IE, the probability of learning is $\frac{pq+(1-p)q^2}{1-q(1-q)}$, while outside is just p. The probability of learning is higher inside the Revelation region if and

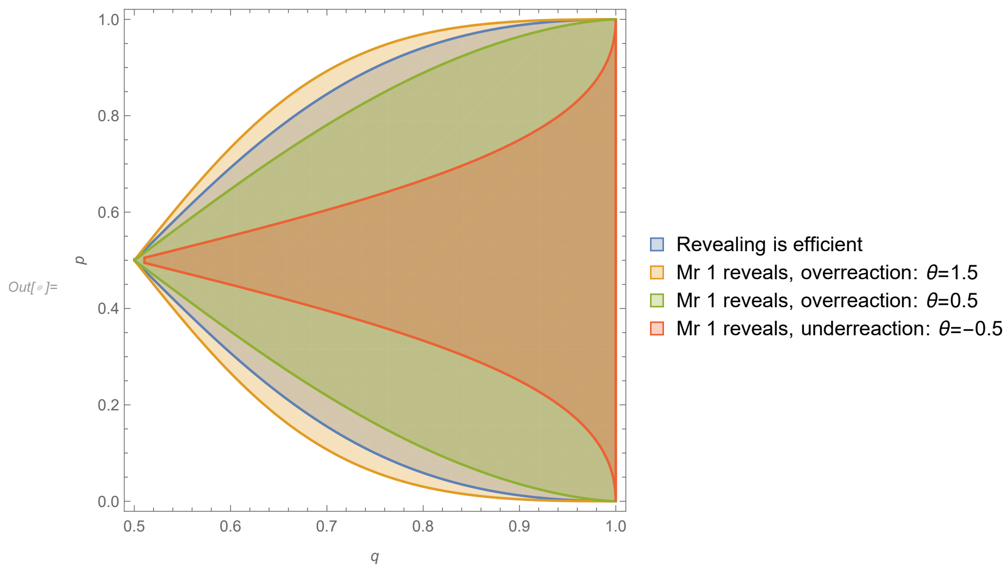

Figure 2: The areas of the parameter space where revealing is efficient, and the area where agent with different degree of distortion reveal. It is clear that under-reaction is always worse than Bayesian, while a moderate over-reaction can be socially better than Bayesian.

only if:

$$p < \frac{pq + (1-p)q^2}{1 - q(1-q)},$$

which is equivalent to:

$$p < \frac{q^2}{(1-q)^2 + q^2}.$$

Now we can rewrite the condition defining the IE region:

$$p < \frac{q^{\theta+1}}{(1-q)^{\theta+1} + q^{\theta+1}}.$$

Depending on θ, $\frac{q^2}{(1-q)^2+q^2}$ can be larger or smaller than $\frac{q^{\theta+1}}{(1-q)^{\theta+1}+q^{\theta+1}}$, with equality for $q = \frac{1}{2}$ $q = 1$, and for all q if $\theta = 1$.

If $\theta < 1$, it means that there is a region outside the IE region (with positive mass, because of the full support assumption on μ) with probability of learning p, strictly smaller than the corresponding probability if it belonged to the IE, hence by increasing θ the probability of learning would increase. If $\theta > 1$, on the contrary, there is a region that belongs to IE where the probability of learning is smaller than p. Hence, the maximum is achieved for $\theta = 1$.

□

4 Simultaneous exchange and market efficiency

The cascades model is very stylized in two major dimensions:

1. signal's likelihood ratios are bounded away from 1 and 0[2]), which means that the informativeness of each signal is limited. (Smith and Sørensen (2000)) identifies this as a necessary condition for the presence of cascades.

2. the model features unit demand and supply.

The goal of this section is to show that the main insights derived in the previous sections are independent of the two features above. Specifically, in this section we explore the implications of diagnostic expectations in a linear quadratic Gaussian model of exchange, similar to the one used by (Bordalo et al. (2020)) or (Grossman and Stiglitz (1980)). Formally, we assume that there is a continuum of agents indexed by $i \in [0, 1]$ and represented with the density f. Each of them can decide her position D_i with respect to the only asset exchanged, where short sales are allowed (D_i can be negative). The price of the asset realized on the market is p and the marginal value of the asset is V. We assume that the utility of an agent i holding D_i units of the asset is:

$$u_i = (V - p)D_i - \frac{1}{2}\gamma D_i^2$$

This functional form, coupled with normality of the variables involved, yields an equilibrium very similar to the one obtained via CARA preferences (as in (Bordalo et al. (2020)), (Grossman and Stiglitz (1980)) and others), but assuming the linear quadratic functional form directly simplifies the welfare analysis, that is what we are interested in.[3]

[2] Or simply *bounded*, in the terminology of (Smith and Sørensen (2000)

[3] This functional form has a risk aversion coefficient of $r_A(D) = \frac{\gamma}{V - p - \gamma D_i}$, that is naturally not constant, but exemplifies how the parameter γ can be taken to represent a measure of the strength of risk aversion (because $r_A(D)$ is increasing in γ for fixed D_i).

Every agent has access to a private signal s_i that has a distribution $\mathcal{N}(V, 1)$ conditional on V. Agents have a common prior on V that is $\mathcal{N}(0, 1)$. Aggregate supply of the asset is random $S \sim \mathcal{N}(0, 1)$, and the realized price is the value p that clears the market, namely p such that $\int D_i(p) f(i) \mathrm{d}i = S$. As for the model in the previous section, we assume agents have diagnostic expectations with parameter θ. This means that they distort their posterior beliefs after observing both the private signal and the price.[4] To solve the model we follow (Bordalo et al. (2020)) in looking for a *diagnostic expectations equilibrium*, namely a pricing function $p(V, S)$, and demand functions $D_i(p, s_i)$, such that:

1. the demand maximizes trader i distorted expected utility given the observation of the private signal s_i and the price p, formally: $D_i(p, s_i) \in \arg\max\{\mathbb{E}^{\theta}[u_i \mid s_i, p(S, V)]\}$. $\mathbb{E}^{\theta}$ is the expectation operator given the diagnostic posterior;

2. the price clears the market: $\int D_i(p(S, V), s_i) f(i) \mathrm{d}i = S$.

In particular, it is crucial that agents in equilibrium use the price as a signal for the fundamental V, and make inference on that to form their posterior.

The model of this section thus departs in two important ways from the discrete cascades model. First, signals are unbounded, and, second, trade volume is not fixed to 1 but is variable. Specifically, the demand/supply of each agent is:

$$D_i = \frac{\mathbb{E}^{\theta}[V \mid I_i] - p}{\gamma} \tag{22}$$

As is standard, the larger the departure of the expected valuation from the price (the asymmetry of information), the more aggressive the trading, both short or long.

The last ingredient we need is welfare. Since in this economy all agents are identical, it is particularly simple and unambiguous.

Definition 4.1. In this economy (ex-ante) *social welfare* is:

$$W = \int \mathbb{E}(u_i) f(i) \mathrm{d}i = \mathbb{E}(u_i)$$

[4]Note that by Theorem 1 for the purposes of the diagnostic updating the order in which agents observe the signal and the price is immaterial.

The last equality follows from the fact that the only heterogeneity in this economy is *after* observing signals, hence the expected utilities of all agents are identical.[5]

First, a calculation yields the following lemma, important for the following: the distorted updating of a normal distribution is still normal.

Lemma 4.1. *If you have a prior $\mathcal{N}(\mu_0, \sigma_0^2)$, and a likelihood $\mathcal{N}(\mu, \sigma^2)$, then the posterior distribution after the diagnostic updating is:*

$$\mathcal{N}\left(\left(\frac{\mu_0}{\sigma_0^2} + (\theta+1)\frac{s}{\sigma^2}\right)\left(\frac{1}{\frac{1}{\sigma_0^2} + \frac{\theta+1}{\sigma^2}}\right), \frac{1}{\frac{1}{\sigma_0^2} + \frac{\theta+1}{\sigma^2}}\right)$$

Using equation (22), the model can be solved through the classic procedure: assume a linear pricing function $p = a + bV + cS$, compute the expectations of each agent, and use the market clearing conditions to find equations for the three parameters a, b, c. The crucial parameter to look at is the informational content of prices, namely b^2/c^2. This is the relevant parameter because conditional on V the price (renormalized, $\frac{p-a}{b}$) has distribution $\mathcal{N}\left(V, \frac{c^2}{b^2}\right)$. Hence, b^2/c^2 is the precision of the price as a signal of V. The next proposition sums up the crucial properties of the solution.

Proposition 4.1. *The precision b^2/c^2 is increasing in θ.*

Society's welfare W is increasing in the precision $\frac{b^2}{c^2}$.

Hence, society's welfare W is increasing *in θ.*

The intuition for the result is the following: over-reaction causes the price to be more informative, increasing the precision $\frac{b^2}{c^2}$. This, in turn, is beneficial for the agents, because as the equation for the demand (22) shows, agents' choices (hence, their utility) crucially depend on their ability to estimate V correctly.

Again, as in the previous model, we find a beneficial effect of over-reaction for society, in a form that is even more extreme that in the cascades example: here the optimal over-reaction would be *infinite*. Note that this happens even if we chose a specification of the utility

[5] In particular, we would obtain the same expression if we started from a more general specification for the utilitarian social welfare: $W = \int \mathbb{E}(u_i)\lambda(i)f(i)\mathrm{d}i$, where $\lambda(i)$ are the Pareto weights.

function that by itself penalizes larger trades (the quadratic term is a cost convex in the trade volume), so, absent the market mechanism, makes *under*-reaction more convenient (because under-reaction generates smaller trade volume). Why is this the case? To understand the mechanics, it is useful to start from a simplified setting in which agents fail to recognize the informational content of the price, and base their forecast only on the private signal, $I_i = \{s_i\}$.

Let us examine equation 22. Market clearing gives

$$p = \int \mathbb{E}[V \mid I_i] f(i) \mathrm{d}i - \gamma S \tag{23}$$

So the price is essentially the average expectation, plus a noise term. This clarifies that, on average, the demand is positive when the agent receives a signal larger than the average $s_i > V$, and negative otherwise. The noise comes from the random supply: agents are not able to distinguish which part of the price is the average expectation and what part is noise. A larger θ means that for the same signal the trader trades more aggressively, and so the expectations part becomes more prominent with respect to the noise. As a consequence, the price is a more precise signal of the fundamental.

To see the mechanism in a more concrete way, consider the following modification of the setting just described. Everything is as in the main model, but agents *fail to understand the informational content of prices.* Technically, definition 4.1 is modified in the following way: agents optimize $\mathbb{E}[u_i \mid s_i]$. In this setting, over-reaction still implies that the price is a more precise signal of the fundamental, but agents do *not* benefit from a the precision of the price signal. As a consequence, over-reaction is not optimal any more, indeed we find that *under-reaction* is optimal instead. The formal result is in the next proposition.

Proposition 4.2. *If agents neglect the informational content of prices, the precision of the price as a signal of the fundamental,* $\frac{b^2}{c^2}$ *is still increasing in* θ*, but society's welfare* W *is* decreasing *in* θ*.*

5 Conclusion

We presented an analysis of the welfare implications of over-reaction, a deviation from Bayesian learning that has recently found a lot of support in macro-financial data. Specifically, we showed through two classic settings used to study asymmetric information in markets, that over-reaction can be beneficial for society because it creates a positive information externality. Over-reacting agents, by having beliefs that move more than bayesians, trade in such a way to reveal more of their private information, either by switching action (in the Glosten-Milgrom discrete setting), or by trading more aggressively (in the linear Gaussian trading volume). A recent research line ask why over-reaction is a so pervasive. [6] Our results suggest that in informationally complex settings, a society or a market with over-reacting agents might have an advantage, thus suggesting an explanation through an evolutionary channel. A further step to understand the consequences of over-reaction is more realistic markets is to include *heterogeneous* beliefs as a wealth of research using survey data documented. This is anevue for future research.

[6] Recent works in this spirit are (**?** and (Afrouzi et al. (2020)).

References

Afrouzi, H., S. Y. Kwon, A. Landier, Y. Ma, and D. Thesmar (2020). Overreaction and working memory. Technical report, National Bureau of Economic Research.

Avery, C. and P. Zemsky (1998). Multidimensional uncertainty and herd behavior in financial markets. *American economic review*, 724–748.

Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics 107*(3), 797–817.

Benjamin, D. J. (2019). Errors in probabilistic reasoning and judgment biases. In *Handbook of Behavioral Economics: Applications and Foundations 1*, Volume 2, pp. 69–186. Elsevier.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy 100*(5), 992–1026.

Bohren, J. A. and D. Hauser (2021). Learning with heterogeneous misspecified models: Characterization and robustness. *Econometrica, forthcoming*.

Bordalo, P., K. Coffman, N. Gennaioli, and A. Shleifer (2016). Stereotypes. *The Quarterly Journal of Economics 131*(4), 1753–1794.

Bordalo, P., N. Gennaioli, S. Y. Kwon, and A. Shleifer (2020). Diagnostic bubbles. *Journal of Financial Economics*.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2018). Overreaction in macroeconomic expectations. Technical report, Working Paper.

Bordalo, P., N. Gennaioli, R. L. Porta, and A. Shleifer (2019). Diagnostic expectations and stock returns. *The Journal of Finance 74*(6), 2839–2874.

Bordalo, P., N. Gennaioli, and A. Shleifer (2018). Diagnostic expectations and credit cycles. *The Journal of Finance 73*(1), 199–227.

Bouchaud, J.-P., P. Krueger, A. Landier, and D. Thesmar (2019). Sticky expectations and the profitability anomaly. *The Journal of Finance 74*(2), 639–674.

Cipriani, M. and A. Guarino (2014). Estimating a structural model of herd behavior in financial markets. *American Economic Review 104*(1), 224–51.

d'Arienzo, D. (2020). Maturity increasing overreaction and bond market puzzles. *Available at SSRN 3733056*.

Decamps, J.-P. and S. Lovo (2006). Informational cascades with endogenous prices: The role of risk aversion. *Journal of Mathematical Economics 42*(1), 109–120.

Frick, M., R. Iijima, and Y. Ishii (2020). Misinterpreting others and the fragility of social learning. *Econometrica 88*(6), 2281–2328.

Glosten, L. R. and P. R. Milgrom (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of financial economics 14*(1), 71–100.

Grossman, S. J. and J. E. Stiglitz (1980). On the impossibility of informationally efficient markets. *The American economic review 70*(3), 393–408.

Lakonishok, J., A. Shleifer, and R. W. Vishny (1992). The impact of institutional trading on stock prices. *Journal of financial economics 32*(1), 23–43.

Maxted, P. (2019). A macro-finance model with sentiment. Technical report, Tech. rep.

Miller, J. W. (2018). A detailed treatment of doob's theorem. *arXiv preprint arXiv:1801.03122*.

Molavi, P., A. Tahbaz-Salehi, and A. Jadbabaie (2018). A theory of non-bayesian social learning. *Econometrica 86*(2), 445–490.

Nofsinger, J. R. and R. W. Sias (1999). Herding and feedback trading by institutional and individual investors. *The Journal of finance 54*(6), 2263–2295.

Park, A. and H. Sabourian (2011). Herding and contrarian behavior in financial markets. *Econometrica 79*(4), 973–1026.

Santosh, S. (2020). Exaggerated likelihoods. *Available at SSRN 3520802*.

Smith, L. and P. Sørensen (2000). Pathological outcomes of observational learning. *Econometrica 68*(2), 371–398.

Tversky, A. and D. Kahneman (1974). Judgment under uncertainty: Heuristics and biases. *science 185*(4157), 1124–1131.

Vives, X. (2017). Endogenous public information and welfare in market games. *The Review of Economic Studies 84*(2), 935–963.

Walther, A. (2020). Prudential policy in an exuberant world.

A Proofs

A.1 Theorem (2

)

Proof. We assumed a buy cascade in the rational cascade starting at time $t = 1$ (the sell cascade case is symmetrical). Namely: $\frac{q}{1-q} > \frac{1-q}{q} > \frac{1-p}{p}$. Also: $\bar{\theta} = 0$ and $\theta > \theta^* = \frac{\frac{p}{1-p}}{\frac{q}{1-q}} + 1$. For $\theta > \theta^*$ in the diagnostic learning case, the first action is perfectly informative of the private signal, because:

$$\left(\frac{q}{1-q}\right)^{1+\theta} > \frac{1-p}{p} > \left(\frac{1-q}{q}\right)^{1+\theta} \tag{24}$$

for $\theta > \theta^*$.

If $s_1 = 1$, then at time $t = 2$:

$$\left(\frac{q}{1-q}\right)^{(1+\theta)s_2} \left(\frac{q}{1-q}\right)^{\theta} > \frac{1-p}{p}. \tag{25}$$

Hence, after $s_1 = 1$ there is buy cascade. If $s_1 = -1$ and $s_2 = -1$ then:

$$\left(\frac{q}{1-q}\right)^{-(1+\theta)-\theta} < \frac{1-p}{p}, \tag{26}$$

hence $H_2 = S$. If instead $s_2 = 1$, then:

$$\left(\frac{q}{1-q}\right)^{(1+\theta)-\theta} > \frac{1-p}{p}, \tag{27}$$

hence $H_2 = B$. Therefore, after $s_1 = -1$, H_2 is perfectly informative. Then, if $s_1 = -1, s_2 = -1$:

$$\left(\frac{q}{1-q}\right)^{(1+\theta)s_3-2\theta} < \frac{1-p}{p}, \tag{28}$$

because we assume θ sufficiently high. Finally, after $s_1 = -1, s_2 = 1$ the trader faces the identical problem of $t = 1$.

The probabilty of a buy cascade given $V = 1$ is therefore:

$$Pr(\text{Buy Cascade}|V = 1) = q + (1 - q)qPr(Pr(\text{Buy Cascade}|V = 1)). \tag{29}$$

Hence: $Pr(\text{Buy Cascade}|V = 1) = \frac{q}{1-(1-q)q}$. □

A.2 Theorem (1)

Proof. **Consistency**

The diagnostic Bayesian operator is defined for $t = 1$ as:

$$\mathcal{BU}_1^\theta(l, p_0)(\omega) = \frac{1}{\int \mathcal{BU}_1(l, p_0) \left(\frac{\mathcal{BU}_1(l,p_0)(\omega')}{p_0(\omega')}\right)^\theta d\omega'} \mathcal{BU}_1(l, p_0) \left(\frac{\mathcal{BU}_1(l, p_0)(\omega)}{p_0(\omega)}\right)^\theta.$$

For $t > 1$ it is defined as:

$$\mathcal{BU}_t^\theta(l, p_0)(\omega) = \frac{1}{\int \mathcal{BU}_t(l, \mathcal{BU}_{t-1}^\theta(l, p_0))(\omega') \left(\frac{\mathcal{BU}_t(l,\mathcal{BU}_{t-1}^\theta(l,p_0))(\omega')}{\mathcal{BU}_{t-1}^\theta(l,p_0)(\omega')}\right)^\theta d\omega'} \mathcal{BU}_t(l, \mathcal{BU}_{t-1}^\theta(l, p_0))(\omega) \Bigg(\frac{\mathcal{BU}_t(l, \mathcal{BU}_{t-1}^\theta(l}{\mathcal{BU}_{t-1}^\theta(l, p}$$

Note that this can be rewritten as:

$$\begin{aligned}
\mathcal{BU}_t^\theta(l, p_0)(\omega) &\propto \mathcal{BU}_t(l, \mathcal{BU}_{t-1}^\theta(l, p_0))(\omega) \left(\frac{\mathcal{BU}_t(l, \mathcal{BU}_{t-1}^\theta(l, p_0))(\omega)}{\mathcal{BU}_{t-1}^\theta(l, p_0)(\omega)}\right)^\theta \\
&\propto \mathcal{BU}_t(l, \mathcal{BU}_{t-1}^\theta(l, p_0))(\omega) \left(\frac{l(X_t|X_1, \ldots, X_{t-1}, \omega)^\theta \mathcal{BU}_{t-1}^\theta(l, p_0)(\omega)}{\mathcal{BU}_{t-1}^\theta(l, p_0)(\omega)}\right)^\theta \\
&\propto l(X_t|X_1, \ldots, X_{t-1}, \omega)^{1+\theta} \mathcal{BU}_{t-1}^\theta(l, p_0)(\omega) \\
&\propto \mathcal{BU}_t(l^{1+\theta}, p_0)(\omega).
\end{aligned}$$

Note that for any $1 < \tau < t$:

$$\mathcal{BU}^{\theta}_{t-\tau}(l, \mathcal{BU}^{\theta}_{\tau}(l, p_0)(\omega))(\omega) = \mathcal{BU}_{t-\tau}(l^{1+\theta}, \mathcal{BU}_{\tau}(l^{1+\theta}, p_0)(\omega))(\omega) = \mathcal{BU}(l^{1+\theta}, p_0)(\omega) = \mathcal{BU}^{\theta}_{t}(l, p_0)(\omega).$$

This says that sequential updating or "one shot" updating are equivalent. Moreover, we have the following simplified expression for the updating:

$$\mathcal{BU}^{\theta}(l, p_0)(\omega) = \frac{l(X|\omega)^{1+\theta} p_0(\omega)}{\int l(X|\omega')^{1+\theta} p_0(\omega') d\omega'}. \tag{30}$$

Asymptotic learning

Theorem 2.4 in Miller (2018) guarantees that the above is true if $\theta = 0$ (Bayesian case). In general:

$$P^{\theta}(\omega \in B | X_1, \ldots, X_t) = \int_B \frac{1}{Z(X_1, \ldots, X_t)} \left(\frac{dP(\cdot|X_1, \ldots, X_t)}{dp_0} \right)^{\theta} dP(\cdot|X_1, \ldots, X_t)$$

is that correct? I don't know. If states are discrete:

$$P^{\theta}(\omega | X_1, \ldots, X_t) = \frac{1}{Z(X_1, \ldots, X_t)} \left(\frac{P(\omega|X_1, \ldots, X_t)}{p_0(\omega)} \right)^{\theta} P(\cdot|X_1, \ldots, X_t)$$

Now for every ω $\left(\frac{P(\omega|X_1, \ldots, X_t)}{dp_0(\omega)} \right)^{\theta} P(\omega|X_1, \ldots, X_t)$ is a continuous function of the Bayesian belief, hence by the continuous mapping theorem, and the result for the Bayesian case, for p_0-almost every ω:

$$\left(\frac{P(\omega|X_1, \ldots, X_t)}{p_0(\omega)} \right)^{\theta} P(\omega|X_1, \ldots, X_t) \to 1_{\omega = \omega^*} p_0(\omega)^{-\theta} \quad l^{\infty}_{\omega^*} - a.s.$$

Moreover:

$$Z(X_1, \ldots, X_t) = \sum_{\omega} \left(\frac{P(\omega|X_1, \ldots, X_t)}{p_0(\omega)} \right)^{\theta} P(\omega|X_1, \ldots, X_t) \to p_0(\omega^*)^{-\theta} \quad l^{\infty}_{\omega^*} - a.s.$$

hence we conclude:

$$\frac{1}{Z(X_1,\dots,X_t)}\left(\frac{P(\omega|X_1,\dots,X_t)}{p_0(\omega)}\right)^{\theta} P(\omega|X_1,\dots,X_t)\to 1_{\omega=\omega^*}\quad l^{\infty}_{\omega^*}-a.s.$$

In the case of continuous signals we have something similar. Assume that the prior and the likelihoods are bounded, so that all the densities are bounded. We know that:

$$F(\omega|X_1,\dots,X_t)\to 1_{\omega\geq\omega^*}\quad l^{\infty}_{\omega^*}-a.s.$$

Now by definition of continuous random variable:

$$F(\omega|X_1,\dots,X_t)=\int_{-\infty}^{\omega} f(\omega'|X_1,\dots,X_t)d\omega'$$

We want to prove that

$$\int_{-\infty}^{\omega} f(\omega'|X_1,\dots,X_t)\left(\frac{f(\omega'|X_1,\dots,X_t)}{f_0(\omega')}\right)^{\theta} d\omega'\to 1_{\omega\geq\omega^*}f_0(\omega^*)^{-\theta}$$

From the convergence of Bayesian distributions we have for any $\delta>0$:

$$\lim_{t\to\infty}\int_{|\omega-\omega^*|>\delta} f(\omega'|X_1,\dots,X_t)d\omega'=0$$

Now we can split the integral above in:

$$\int_{|\omega'-\omega^*|>\delta,\omega'\leq\omega} f(\omega'|X_1,\dots,X_t)\left(\frac{f(\omega'|X_1,\dots,X_t)}{f_0(\omega')}\right)^{\theta} d\omega'+\int_{|\omega'-\omega^*|\leq\delta,\omega'\leq\omega} f(\omega'|X_1,\dots,X_t)\left(\frac{f(\omega'|X_1,\dots,X_t}{f_0(\omega')}\right.$$

The first integral is smaller than $\int_{|\omega-\omega^*|>\delta} f(\omega'|X_1,\dots,X_t)d\omega'\max_{|\omega'-\omega^*|>\delta} f_0\to 0$ by the property above, while the second is equal to, by the mean value theorem there exists a $\overline{\omega}$ such that:

$$f_0(\overline{\omega})^{-\theta}\int_{|\omega'-\omega^*|\leq\delta,\omega'\leq\omega} f(\omega'|X_1,\dots,X_t)^{\theta+1}d\omega'$$

where $|\omega^*-\overline{\omega}|\leq\delta$.

Now, if $\omega<\omega^*$ also this second integral goes to zero, otherwise it goes to 1, and by continuity of the prior we conclude:

$$\int_{-\infty}^{\omega} f(\omega'|X_1,\ldots,X_t)\left(\frac{f(\omega'|X_1,\ldots,X_t)}{f_0(\omega')}\right)^{\theta} d\omega' \to 1_{\omega\geq\omega^*} f_0(\omega^*)^{-\theta}$$

repeating the reasoning for the normalization gets rid of the prior and concludes the proof.

□

A.3 Proof of Proposition 3.2

Mr 1 will play action $a_1 = 1$ after observing signal $s_1 = 1$ if and only if his posterior is higher than τ, that is:

$$\frac{pq^{1+\theta}}{pq^{1+\theta}+(1-p)(1-q)^{1+\theta}} \geq \tau$$

$$\iff 1+\frac{1-p}{p}\left(\frac{1-q}{q}\right)^{1+\theta} \leq \frac{1}{\tau}$$

$$\iff (1+\theta)\log\left(\frac{1-q}{q}\right) \leq \log\left(\frac{1}{\tau}-1\right)\frac{p}{1-p}$$

$$\iff \theta+1 \geq \frac{\ln\left(\frac{p}{1-p}\left(\frac{1}{\tau}-1\right)\right)}{\ln\frac{1-q}{q}} = \frac{-\ln\left(\frac{p}{1-p}\right)+\ln\left(\frac{\tau}{1-\tau}\right)}{\ln\frac{q}{1-q}}.$$

In the last line we used $q > \frac{1}{2}$ to change the inequality sign. This condition is always true if $p \geq \tau$, given that $\theta > -1$. Its interpretation is that if Mr 1 is ex ante indifferent or in favor of alternative 1, then if he observes signal $s_1 = 1$, he plays $a_1 = 1$ for any value of θ. On the contrary if $p < \tau$, meaning that the agent is ex ante in favor of alternative 0, when he observes signal $s_1 = 1$ he might or might not play action $a_1 = 1$, depending on the parameter values. The condition above says that the bigger θ is, the bigger the set of parameters under which Mr1 revises the prior and plays action $a_1 = 1$.

Similarly, after seeing $s_1 = 0$, Mr 1 will play action $a_1 = 1$ if and only if:

$$\frac{p(1-q)^{1+\theta}}{p(1-q)^{1+\theta}+(1-p)q^{1+\theta}} \geq \tau$$

$$\iff \theta+1 \leq \frac{\ln\left(\frac{p}{1-p}\left(\frac{1}{\tau}-1\right)\right)}{\ln\frac{q}{1-q}} = \frac{\ln\left(\frac{p}{1-p}\right)-\ln\left(\frac{\tau}{1-\tau}\right)}{\ln\frac{q}{1-q}}.$$

This is never the case never if $p < \tau$, which means that if the agent is in favor of alternative 0 and then he sees the signal $s_1 = 0$, he never revises his opinion. On the contrary, depending on θ, the opposite case may be true. Call the above condition 2. Note that the space of parameter such that condition two is violated increases with θ.
Summing up: if the agent sees $s_1 = 1$, then he plays $a_1 = 1$ the if either $\tau \leq p$ or $\tau > p$ and condition 1 is satisfied. If the agent sees $s_1 = 0$, then he plays $a_1 = 0$ if $\tau > p$ or $\tau \leq p$ and condition 2 is violated.

The behavior of Mr 2 will depend on which conditions are satisfied. Consider the informationally efficient region IE (for agent 1) defined as:

$$IE = \{(p, \tau) \in [0,1] \times \mathbb{R}_+ | \, (\tau \leq p \text{ and condition 2 is true}) \, \text{or} \, (\tau > p \text{ and condition 1 is true })\}.$$

If the parameters lie inside IE, then Mr 2 can perfectly infer Mr1 signal by observing his action, since $a_1 = s_1$. Thus Mr 2 effectively observes two signals.
Consider first the case $p > \tau$, namely the prior is in favor of 1. In this case if $s_1 = 1$ than Mr2 will do his Bayesian updating, leading him to play action $a_2 = 1$ regardless of his signal. Similarly for subsequent agents: a cascade therefore starts on state 1 in this case. If instead $a_1 = s_1 = 0$, then if $s_2 = 0$ Mr2 will do his Bayesian updating, leading him to play action $a_2 = 0$ regardless of his signal. Similarly for subsequent agents: a cascade therefore starts on state 0 in this case. Finally if $a_1 = s_1 = 0$ and $s_2 = 1$, then Mr2 will do his Bayesian updating, leading him to play action $a_2 = 1$. However in this case a cascade does not start immediately, as Mr 3 faces the same problem of Mr1. Here the intuition is straightforward: opposite signals s_1 and s_2 cancel out, and therefore $Mr2$ only relies on his prior belief. The case $p < \tau$ is symmetric.
Resuming the dynamics is characterized as follows:

- if $p > \tau$ then the probability of learning is:

$$Pr(a_\infty = \omega) = pPr(a_\infty = 1) + (1-p)Pr(a_\infty = 0) = \big(pq + (1-p)q^2\big) \left(\sum_{i=0}^{\infty} (q(1-q))^i \right);$$

- if $p < \tau$ then the probability of learning is:

$$Pr(a_\infty = \omega) = pPr(a_\infty = 1)+(1-p)Pr(a_\infty = 0) = \left(pq^2 + (1-p)q\right)\left(\sum_{i=0}^{\infty}(q(1-q))^i\right);$$

Resuming, if $s_1 = s_2$, then $a_2 = s_2$; if instead $s_1 \neq s_2$ then M2 2 will stick to his prior belief. In the former case Mr 3 will also play $a_3 = s_2$ regardless of his signal (if $s_3 = s_2 = s_1$ this is true since $IE_1 \subseteq IE_2 \subseteq IE_3$; if $s_3 \neq s_2 = s_1$ then Mr3 problem is the same problem of Mr1, therefore $a_3 = a_3$); in the latter case Mr3 problem is the same problem of Mr 1, therefore $a_3 = s_3$.
Therefore, if the parameters lie IE_1, then i

plays 1 if $p \geq \tau$, and 0 viceversa. This means that if $p \geq \tau$ and the first agent revealed his signal to be 1, then the second agent will always play 1, and this will not be informative for Mr 3, which will act as if he observed only the signal of the first agent. On the contrary, if the first agent revealed his signal to be 0 and still $p \geq \tau$, the second agent reveals his signal, and Mr 3 updates consequently. Hence, if the conditions on θ for Mr 1 are satisfied, the first agent reveals and the second follows if observes the same, and if observes a different signal it depends on the prior, as should be. If the conditions are satisfied for the first but not the second agent, it means that the first agent actually does not reveal information, hence the second agent actually behaves as the first, and it means that he will not reveal anything either, and we have the applicable cascade (because all subsequent agents will follow).

f Mr 1 plays 1 regardless of the signal s_1 observed, then Mr 2 has no updating to do, and will act as if she were the first of the line. This happens with probability $1 - q$.

Hence, if agents are all homogeneous, there is a trivial cascade on 1, and the probability of learning is p.

Mr 3 if he observes 2 identical zeros will ignore his private signals, and we have a cascade on 0 (the conditions on θ are trivially satisfied). If he observes 3 different signals, will follow the most frequent. Anyway, the first 2 signals are sufficient to determine which cascade we have.

A.4 Proof of Proposition 4.1

Taking first order conditions we obtain that the demand of agent i is:

$$D_i = \frac{E_i^\theta(V \mid p, s_i) - p}{\gamma}$$

Hence, by market clearing, the equilibrium price is:

$$p = \int E_i^\theta(V \mid p, s_i) f(i) di - \gamma S$$

Now, let us assume that the price follows the pricing function $p = a + bV + cS$. Then, the price is a signal distributed as $(p-a)/b \mid V \sim \mathcal{N}(V, c^2/b^2)$. Using the updating rule for diagnostic expectations, we find:

$$V \mid s_i, p \sim \mathcal{N}\left(\left(\frac{\frac{(\theta+1)X_i}{\theta+2}}{\frac{1}{\theta+2}} + (\theta+1)\frac{p-a}{b}\frac{b^2}{c^2}\right)\frac{1}{\theta+2+\frac{(\theta+1)b^2}{c^2}}, \frac{1}{\theta+2+(\theta+1)\frac{b^2}{c^2}}\right)$$

so that, in equilibrium:

$$p = \int \frac{(\theta+1)s_i + (\theta+1)\frac{p-a}{b}\frac{b^2}{c^2}}{\theta+2+\frac{(\theta+1)b^2}{c^2}} f(i) di - \gamma S$$

by the law of large numbers $\int X_i f(i) di = V$ and so:

$$p = \frac{\mu_0 + (\theta+1)\left(V + \frac{p-a}{b}\frac{b^2}{c^2}\right)}{\theta+2+\frac{(\theta+1)b^2}{c^2}} - \gamma S$$

Solving for p:

$$p\left(1 - (\theta+1)\frac{\frac{1}{b}\frac{b^2}{c^2}}{\theta+2+\frac{(\theta+1)b^2}{c^2}}\right) = \frac{(\theta+1)\left(V - \frac{a}{b}\frac{b^2}{c^2}\right)}{\theta+2+\frac{(\theta+1)b^2}{c^2}} - \gamma S$$

so we get the equations for the model parameters:

$$a = \frac{-(\theta+1)\frac{a}{b}\frac{b^2}{c^2}}{\theta+2+\frac{(\theta+1)b^2}{c^2}} \left(\frac{\theta+2+(\theta+1)\left(1-\frac{1}{b}\right)\frac{b^2}{c^2}}{\theta+2+\frac{(\theta+1)b^2}{c^2}} \right)^{-1}$$

from which it follows $a = 0$;

$$b = \frac{(\theta+1)}{\theta+2+\frac{(\theta+1)b^2}{c^2}} \left(\frac{\theta+2+(\theta+1)\left(1-\frac{1}{b}\right)\frac{b^2}{c^2}}{\theta+2+\frac{(\theta+1)b^2}{c^2}} \right)^{-1} = \frac{\theta+1}{\theta+2+(\theta+1)\left(1-\frac{1}{b}\right)\frac{b^2}{c^2}}$$

$$c = -\gamma \left(\frac{\theta+2+(\theta+1)\left(1-\frac{1}{b}\right)\frac{b^2}{c^2}}{\theta+2+\frac{(\theta+1)b^2}{c^2}} \right)^{-1}$$

b and c can be further solved as:

$$b = \frac{(\theta+1)b}{(\theta+2)b+(\theta+1)(b-1)\frac{b^2}{c^2}}$$

$$(\theta+2)b+(\theta+1)(b-1)\frac{b^2}{c^2} = \theta+1$$

$$b = \frac{(\theta+1)\left(1+\frac{b^2}{c^2}\right)}{\theta+2+(\theta+1)\frac{b^2}{c^2}}$$

and

$$b = \frac{(\theta+1)\left(1+\frac{b^2}{c^2}\right)}{\theta+2+(\theta+1)\frac{b^2}{c^2}} = \frac{\left(1+\frac{b^2}{c^2}\right)}{\frac{1}{\theta+1}+1+\frac{b^2}{c^2}} = \left(1+\frac{b^2}{c^2}\right)\gamma\frac{b}{c}$$

so that:

$$c = \left(1+\frac{b^2}{c^2}\right)\gamma$$

It will be useful to express everything in terms of the precision of the price signal, that is:

$$\frac{b^2}{c^2} = \frac{1}{\gamma^2\left(\frac{1}{\theta+1}+\left(1+\frac{b^2}{c^2}\right)\right)^2}$$

and by the implicit function theorem can be seen that it is *increasing* in θ, which is our first thesis.

Finally, the price is:

$$p = \left(1 + \frac{b^2}{c^2}\right) \gamma \left(V \frac{b}{c} - S\right)$$

and so, the demand:

$$D_i = E_i^\theta V - p = \frac{(\theta+1)s_i + (\theta+1)\frac{p}{b}\frac{b^2}{c^2}}{\theta + 2 + \frac{(\theta+1)b^2}{c^2}} - p =$$

$$= \frac{(\theta+1)s_i}{\theta + 2 + \frac{(\theta+1)b^2}{c^2}} - p \frac{\theta + 2 + (\theta+1)\left(1 - \frac{1}{b}\right)\frac{b^2}{c^2}}{\theta + 2 + \frac{(\theta+1)b^2}{c^2}}$$

$$= -\gamma \frac{b}{c} s_i - \left(1 - \frac{1}{b} \frac{(\theta+1)\frac{b^2}{c^2}}{\theta + 2 + (\theta+1)\frac{b^2}{c^2}}\right) p = -\gamma \frac{b}{c} s_i - \left(1 - \frac{1}{\gamma \frac{b}{c}\left(1 + \frac{b^2}{c^2}\right)} \gamma \frac{b^3}{c^3}\right) p$$

$$= -\gamma \frac{b}{c} s_i - \frac{1}{\left(1 + \frac{b^2}{c^2}\right)} p$$

Welfare

We now have all the elements to calculate total welfare.

$$\frac{1}{\gamma}\mathbb{E}[(V-p)(E_i^\theta V - p) - \frac{1}{2}(E_i^\theta V - p)^2] = \frac{1}{\gamma}\mathbb{E}[\mathbb{E}[(V-p)(E_i^\theta V - p) - \frac{1}{2}(E_i^\theta V - p)^2 \mid V]]$$

and

$$D_i = E_i^\theta V - p = \frac{(\theta+1)s_i + (\theta+1)\frac{p}{b}\frac{b^2}{c^2}}{\theta + 2 + \frac{(\theta+1)b^2}{c^2}} - p =$$

$$= \frac{(\theta+1)s_i}{\theta + 2 + \frac{(\theta+1)b^2}{c^2}} - p \frac{\theta + 2 + (\theta+1)\left(1 - \frac{1}{b}\right)\frac{b^2}{c^2}}{\theta + 2 + \frac{(\theta+1)b^2}{c^2}}$$

$$= \frac{1}{\left(\theta + 2 + \frac{(\theta+1)b^2}{c^2}\right)} \left((\theta+1)X_i - (\theta+1)V - (\theta+1)\frac{c}{b}S\right)$$

$$= \gamma \left(-\frac{b}{c}(X_i - V) + S\right)$$

while

$$V - p = (1-b)V - cS$$

so

$$\mathbb{E}[(V-p)(E_i^\theta V - p) \mid V] = E[((1-b)V - cS)\gamma(-\frac{b}{c}(s_i - V) + S) \mid V] =$$

$$\mathbb{E}[\mathbb{E}[((1-b)V-cS)\gamma(-\frac{b}{c}(s_i-V)+S) \mid V,S] \mid V] = \mathbb{E}[((1-b)V-cS)\gamma\mathbb{E}[(-\frac{b}{c}(s_i-V)+S) \mid V,S] \mid V]$$

$$\mathbb{E}[((1-b)V-cS)\gamma\mathbb{E}[S \mid V,S] \mid V] = \mathbb{E}[((1-b)V-cS)\gamma S \mid V] = -\gamma c$$

and:

$$\mathbb{E}[(E_i^\theta V-p)^2 \mid V] = \gamma^2\mathbb{E}[(-\frac{b}{c}(s_i-V)+S)^2 \mid V] = \gamma^2\left(\frac{b^2}{c^2}+1\right)$$

so in total:

$$W = \gamma\left(-c-\frac{1}{2}\gamma\left(\frac{b^2}{c^2}+1\right)\right) = \frac{1}{2}\gamma\left(\frac{b^2}{c^2}+1\right)$$

that is increasing in θ.

A.5 Proof of Proposition 4.2

After having observed the signal s_i, the posterior of agent i is:

$$V \mid s_i \sim \mathcal{N}\left(((\theta+1)s_i)\left(\frac{1}{\theta+2}\right), \frac{1}{\theta+2}\right)$$

where we maintain that the agent does not make any inference on V from the value of the price p. So: $D_i = \frac{\left(\frac{(\theta+1)s_i}{\theta+2}-p\right)}{\gamma} = \frac{\frac{(\theta+1)s_i}{\theta+2}-p}{\gamma}$, and the equilibrium price is:

$$p = \frac{\int((\theta+1)s_i)f(i)\mathrm{d}i}{\theta+2} - \gamma S = \frac{(\theta+1)V}{\theta+2} - \gamma S$$

by the law of large numbers.

Welfare

The total welfare is:

$$\frac{1}{\gamma}\mathbb{E}[(V-p)(E_i^\theta V-p)-\frac{1}{2}(E_i^\theta V-p)^2] = \frac{1}{\gamma}\mathbb{E}[\mathbb{E}[(V-p)(E_i^\theta V-p)-\frac{1}{2}(E_i^\theta V-p)^2 \mid V]]$$

and

$$E_i^\theta V-p = \frac{(\theta+1)s_i}{\theta+2} - \frac{(\theta+1)V}{\theta+2} + \gamma S = \frac{(\theta+1)(s_i-V)}{\theta+2} + \gamma S$$

$$V-p = V - \frac{(\theta+1)V}{\theta+2} + \gamma S = \frac{V}{\theta+2} + \gamma S$$

and

$$\mathbb{E}[(V-p)(E_i^\theta V - p) \mid V] = \mathbb{E}[\left(\frac{V}{\theta+2} + \gamma S\right)\left(\frac{(\theta+1)(s_i - V)}{\theta+2} + \gamma S\right) \mid V]$$

$$= \mathbb{E}[\left(\frac{V}{\theta+2} + \gamma S\right)(\gamma S) \mid V] = \gamma^2$$

and

$$\mathbb{E}[(E_i^\theta V - p)^2 \mid V] = \frac{(\theta+1)^2}{(\theta+2)^2} + \gamma^2$$

so the total is:

$$\frac{1}{2}\left(\gamma^2 - \frac{(\theta+1)^2}{(\theta+2)^2}\right)$$

that is decreasing in θ.