Disagreeing forever: a testable model with non-vanishing belief heterogeneity

Arthur Beddock*

March 30, 2021

First draft

Abstract

I model an economy populated by a large number of investors who persistently differ in their beliefs. Building on the seminal work of Lindahl (1939), I allow the agents to revise their consumption plans in a time-inconsistent way when new information arrives, leading to non-vanishing heterogeneity and empirically testable results. The model is fully equivalent to, and can thus be described as, an overlapping heterogeneous generations model. I find that the equilibrium stock mean return and volatility increase with belief dispersion. Using analyst forecasts from the IBES database, I show evidence that these positive relations are empirically verified when considering the overall market.

Keywords: Heterogeneous beliefs, Continuum of agents, Consumption plan revision, Asset pricing

JEL classification: D53 - D90 - G12 - G40

*Université Paris Dauphine - PSL, Department of Finance, Place du Maréchal de Lattre de Tassigny, 75016 Paris, France, and Tilburg University, Department of Finance, PO Box 90153, 5000 LE Tilburg, The Netherlands. Email: `arthur.beddock@dauphine.psl.eu`

1 Introduction

It is well-established that there is a great belief heterogeneity among market participants. Using a recent survey administered to a panel of retail investors, Giglio et al. (2020) for instance show that beliefs are characterized by a large and persistent individual heterogeneity, and that investors are likely to exhibit a willingness to "agree to disagree". While early studies argued that traders with biased beliefs could be neglected, the recent literature (discussed below) has shown these views to be wrong, and has extensively studied, both from empirical and theoretical perspectives, the implications of such belief dispersion on market characteristics. In particular, Jouini and Napp (2011) demonstrate that heterogeneity is important per se, as an economy with biased investors who are rational on average differs markedly from an economy with rational investors only. Thus, belief heterogeneity matters, and theoretical studies of its impacts need to be empirically tested in a careful manner.

In this paper, I develop and empirically test a model that is able to match empirically observed market features. I build on Atmaz and Basak (2018), and develop a dynamic general equilibrium model evolving in continuous time in an infinite horizon economy populated by a continuum of constant relative risk aversion (CRRA) investors who agree to disagree and can revise their actions in a time-inconsistent way. This latter property is discussed in more detail later, and is of first importance as it allows me, unlike existing theoretical studies, to maintain a persistent heterogeneity among investors, and therefore to empirically test the model over long periods. Thus, I first solve for the model equilibrium and derive its characteristics, before testing it using data on analysts' forecasts from the Institutional Brokers Estimate System (IBES) Unadjusted Summary database.

An important feature of my model is the use of a continuum of agents, which has two main advantages. First, it allows me to consider an unbounded investor type space, or, stated differently, to take into account every possible existing belief. As discussed in Atmaz and Basak (2018), it implies that belief dispersion does not vanish in relatively extreme states. In a model with bounded belief biases, the most biased agents (towards optimism or pessimism) would eventually control almost all of the wealth in the economy in the most extreme states (i.e. the very good or very bad ones), and belief heterogeneity would disappear. Combined with the agents' ability to revise their actions as time goes by, the choice of a continuum thus guarantees that belief heterogeneity persists for all dates and all states of the world. The second advantage of considering an infinite number of investors is that I can use a statistical distribution to describe their wealth shares. Following, e.g., Cvitanic and Malamud (2011), I assume that it follows a Gaussian distribution. Thus, I only need to estimate two parameters, namely the average belief bias and the belief dispersion, in order to describe agents' beliefs, which eases the model empirical assessment. Moreover, it allows me to easily disentangle between the impacts of the first and second moments. This is interesting because, as underlined by Brandon and Wang (2020), most models focus on the effects of the latter.

More importantly, I allow the investors to change their plans as time goes by. In fact, my main methodological contribution is to adapt the well-known and widely used Arrow-Debreu framework to allow the investors to revise their plans in a time-inconsistent way as new information comes in, leading to stationary and testable results. A model endowed with this assumption is actually fully equivalent to an overlapping generations model, and my model can thus alternatively be described as such. More precisely, the time-inconsistency of the investors' consumption plans could equally result from the fact that the plans are successively made by different generations of heterogeneous investors. Under this alternative interpretation—informally detailed in the theoretical part prior to the formal definition of the model—each generation of heterogeneous investors would have the opportunity to consume part of its wealth and to bequest the remaining part to the subsequent generation. For ease of exposition, I nevertheless mostly focus on the equivalent model definition that builds on Lindahl (1939)'s theory, and explain how the two descriptions relate.

Allowing for investors' plan revision as time goes by echoes the seminal work of Lindahl (1939), who observes that "the plans of the economic subjects at any given point of time are neither fully consistent with one another nor with the external conditions, and therefore they must be successively revised."[1] This is also in line with the more general temporary general equilibrium theory of Grandmont (1977, 2008). Thus, I assume continuous trading over time and continuous revisions of the agents' plans in response to the arrival of new information on the state of the world. Note that, at each date t, I further posit that the agents consider a finite horizon economy which is supposed to end at date $t + T$, and that they make their plans accordingly before revising them within a similar framework at the subsequent date $t + dt$. Modeling an economy à la Lindahl (1939) thus leads to a infinite continuum of temporary equilibria with sliding horizon instead of an infinite rational expectation equilibrium. Under the alternative overlapping generations model description, this corresponds to a model with an infinite number of generations that all have a life of length T and where Generation-t gives birth to a similar new generation at time $t + dt$.

Importantly, and in line with what is empirically observed, this continuous plans' revision feature ensures that the belief heterogeneity is persistent as consumption plan revisions prevent every investor to go extinct. This is also in sharp contrast with the Arrow-Debreu framework, which relies on the assumption that all plans prevailing at the starting date of the economy are dynamically time consistent, implying an unrealistically high level of agent rationality. Additionally, the model is therefore well-suited to be tested empirically

[1] Erik Lindahl, *Studies in the theory of money and capital*, 1939: p.38. Lindahl also indicates a general procedure to construct a solution consistent with such successive plans revisions:

> Starting from the plans and the external conditions valid at the initial point of time, we have first to deduce the development that will be the result of these data for a certain period forward during which no relevant changes in the plans are assumed to occur. Next we have to investigate how far the development during this first period—involving as it must various surprising for the economic subjects—will force them to revise their plans of action for the future, the principles for such a revision being assumed to be included in the data of the problem. And since on this basis the development during the second period is determined in the same manner as before, fresh deductions must be made concerning the plans for the third period, and so on.

over long periods.

Let me now turn to the model's implications, and see how it relates to the existing literature. Because the modeling methodologies share similarities, most of the theoretical results are similar to those in Atmaz and Basak (2018) evaluated at $t = 0$. However, let me emphasize that, unlike their finite horizon model where the heterogeneity effects progressively fade away, I derive stationary results. This feature is crucial as it is now well-documented that belief heterogeneity is empirically persistent through time (see, e.g., Meeuwis et al., 2019, Giglio et al., 2020, Das et al., 2020).

Looking at the stock price, I infer that it depends positively on the average belief bias, which is in line with the studies of Jouini and Napp (2007) and Kurz and Motolese (2011). I also derive that the impact of belief dispersion is positive for sufficiently good states of the world and negative for sufficiently bad ones. Interestingly, the sign of this impact only depends on t through the current state of the world W_t, and not through the remaining time before the economy ending date as this is the case in a finite horizon setting. Finally, similarly to Atmaz and Basak (2018), I find a convex relation between the stock price and the cashflow news. This price convexity implies that the stock price reacts more to good news than to bad ones, and that the stock price reaction to any type of news is stronger in relatively good states. Basu (1997) and Nagel (2005) provide empirical evidence for the first prediction, and, consistent with the second one, Conrad et al. (2002) show that the market responds more strongly to bad news in good times than in bad times. Other theoretical studies derive this convex relation in a model with incomplete common information (Veronesi, 1999), or assuming short-sale constraints (Xu, 2007).

I also study the relation between belief heterogeneity and the stock mean return, and, unlike Atmaz and Basak (2018), I observe that the higher the heterogeneity, the higher the expected return is. The positive relation that I document is in line with the conjecture of Williams (1977) that more dispersion of opinion represents more risk, and therefore that agents should be more compensated for holding a riskier asset. Banerjee and Kremer (2010) confirm this predicted positive relation in a dynamic model in which investors disagree on the interpretation of public information, and Buraschi and Jiltsov (2006) derive a similar result linking heterogeneity in beliefs to option open interest. This view is also empirically supported (see, e.g., Anderson et al., 2005, Doukas et al., 2006, Banerjee, 2011). In addition, I obtain that the stock mean return unconditionally decreases with risk aversion. This is because, in a heterogeneous economy, more risk averse investors speculate less aggressively, and thus earn lower returns.

I further derive that the stock volatility monotonically increases with belief dispersion, and is persistently higher than the production process volatility. As stated in Atmaz and Basak (2018), this is because higher fluctuations in the average belief bias translate to additional stock price fluctuations and therefore increase stock volatility. This monotonic positive relation between belief dispersion and stock volatility is well-documented both theoretically (see, e.g., Shalen, 1993 in a two-period rational expectations model, Scheinkman and Xiong, 2003 in a model with short sale constraints, Buraschi and Jiltsov, 2006 in a model with rational agents with incomplete and heterogeneous information, Andrei et al., 2019 in a model with disagreement on the length of business cycles) and empirically (see, e.g., Ajinkya and Gift, 1985, Anderson et al., 2005, Banerjee, 2011). I complement these findings by deriving a stationary formula where the heterogeneity effects on volatility remain persistent over time.

Turning to the empirical part of the paper, I provide an analysis of the above mentioned positive relations by running ordinary least squares (OLS) regressions. More precisely, I focus on studying the empirical ability of belief dispersion to predict future market returns and volatility over time using variables sampled at the quarterly frequency. Rolling window regressions complete the analysis and allow to more carefully study the time evolution of the belief dispersion predictive ability. As indicated before, the—empirically observed—persistence of my theoretical predictions is key as they can thus be estimated over long periods.

Following the literature, I use analyst monthly forecasts of the earnings-per-share (EPS) long-term growth rate (LTG) of individual stocks from the IBES Unadjusted Summary database from January 1982 to December 2019 as a proxy for investors' beliefs. Building on Yu (2011), I then aggregate them over time and across assets to obtain quarterly market belief dispersion data. I further compute the market annualized quarterly (raw) returns from data on individual stock prices from the Center for Research in Security Prices (CRSP) database, defining the market each month as the portfolio containing all individual stocks for which at least two monthly EPS LTG forecasts are provided in the IBES database. Similarly, quarterly annualized market volatility data are constructed from daily returns of this portfolio. Therefore, the same assets are used for the construction of the three variables of interest, which allows me to more precisely capture the link between the market characteristics and the investor beliefs.

Overall, the empirical results point towards the approval of the model-implied positive effects of belief dispersion on the market returns and volatility. The rolling window regressions further show that the sign and intensity of the belief dispersion impact varies through time, which might explain why the results are not statistically significant for all specifications considered.

This paper adds on the belief heterogeneity literature. In particular, it contributes to the large debate on the impact of belief dispersion on the stock mean return. While I derive a positive relation, a strand of the literature, based on the seminal work of Miller (1977), documents a negative one.[2] This negative relation critically depends on the presence of market frictions. For example, Chen et al. (2002) obtain this result by developing a model with differences of opinion and short-sales constraints. On the empirical side, Diether et al. (2002) report that high dispersion stocks earn lower returns. Interestingly, Doukas et al. (2006) replicate their results, and find that the relation becomes positive when controlling for uncertainty in analysts' earnings forecasts. Other studies derive mixed results or no relation. In particular, Buraschi et al. (2014) find that the relation is ambiguous and leverage dependent: it is positive and significant for high leverage firms, but can turn negative and non significant for moderately leveraged firms. Atmaz and Basak (2018)

[2] In a model with short-sale constraints and differences of opinion, Miller (1977) argues that the stock is overpriced as it reflects the view of the optimistic agents. In fact, because of the short-sale constraints, pessimistic agents stay out of the market. The higher the differences of opinion, the higher this effect, and therefore the higher the stock overpricing, resulting in lower subsequent returns.

theoretically derive that higher dispersion leads to higher returns when the view on the stock is sufficiently pessimistic, and to lower returns when the view is sufficiently optimistic. Finally, Avramov et al. (2009) find that financial distress drives the reported negative dispersion effect, and show that it is a facet of non-investment grade firms which account for less than 5% of the total market capitalization, and that the effect is virtually non existent otherwise.

The paper is organized as follows. Section 2 contains a brief description of the alternative overlapping generations model. It also presents the main model, its equilibrium characteristics, and derives the theoretical results relative to the stock price, its mean return, and its volatility. Using empirical data, I test the model in Section 3. Section 4 concludes. All proofs are reported in Appendix A, and Appendix B shows that the two alternative models are equivalent.

2 Theoretical part

This section shows how the model equilibrium and representative agent are computed based on Lindahl (1939)'s theory, and derive the main theoretical implications. Such a description of the model has the advantage to ease its formal definition and resolution. Before doing so, let me describe informally a fully equivalent model that relies on overlapping heterogeneous generations, which allows to get more intuition, and yields similar implications.

2.1 An overlapping heterogeneous generations model

Consider an infinite horizon economy evolving in continuous time populated by a large number of overlapping generations of investors with heterogeneous beliefs who maximize their utility. The mechanism of such a model is the following.

Each generation has a life of length T, and each investor of a given generation has her own beliefs regarding the future of the economy. These beliefs—more formally defined in Section 2.2—are indexed by a coefficient δ. At any time t, Generation-t is born, and Agent-

δ of Generation-t gives birth to Agent-δ of Generation-$t + dt$ at time $t + dt$. These agents are thus part of Family-δ. From a generation to another, the investors of a same family therefore keep the same beliefs which allows to maintain a persistent belief heterogeneity in the economy as time goes by.

There is a production process $\tilde{y}$ in the economy that follows a geometric Brownian motion, and each agent of a given generation is endowed with a fraction of this production process. The heterogeneous investors can then trade at an endogenously determined price depending on their beliefs in order to maximize their utility. In equilibrium, at time $t+T$, each agent of Generation-t consumes a part of their endowment, bequests the remaining part to Generation-$t + dt$, and dies. Thus, at time $t + T$, Generation-$t + dt$ inherits part of the endowment of Generation-t. At time $t + T + dt$, Generation-$t + dt$ consumes a part of the bequested endowment, bequests the remaining part to the next generation, and dies. The same reasoning applies to all subsequent generations and thus results in a continuum of consumption and bequest plans made by all successive generations. Because each generation is born at a different time and in a different state of the world, the successive plans of the members of a same family are not necessarily time consistent.

This overlapping heterogeneous generations model is equivalent to the model based on Lindahl (1939)'s theory presented below where there is no effective consumption. In fact, Appendix B shows that both lead to the exact same implications in the case where each generation only bequests to the next one. When effective consumption is allowed, there is only a small adjustment to be made regarding the drift of the production process $\tilde{y}$ to obtain similar results in both models. Without loss of generality, I thus focus on the former case in the remainder of the analysis.

2.2 Equilibrium and representative agent

Let me now more formally define the model based on Lindahl (1939)'s theory, and solve for its equilibrium before deriving its properties.

I define a pure-exchange security market economy à la Lindahl (1939) evolving in continuous time with an infinite horizon. The economy is populated by heterogeneous agents who maximize their expected utility from future consumption, and revise their plans as time goes by.

Uncertainty is modeled by a filtered probability space $(\Omega, F, (F_t), \mathbb{P})$, where Ω is the set of states of nature, F is the σ-algebra of observable events, (F_t) describes how information is revealed through time, and $\mathbb{P}$ is the true probability measure giving the likelihood of occurrence of the different events in F. I assume that there is a single source of risk, modeled by a $((F_t), \mathbb{P})$-Brownian motion W.

As formally explained later, let y denote the expected production process in the economy, and assume that, under the true probability measure $\mathbb{P}$, it follows a geometric Brownian motion with a drift μ and a volatility σ. This expected production process is comparable to the process $\tilde{y}$ given in Section 2.1.

There is a large number of heterogeneous investors in the economy. More precisely, I assume that the economy is populated by a continuum of infinitely lived agents with heterogeneous beliefs. All agents have standard CRRA preferences, characterized by $u(x) = \frac{x^{1-\gamma}}{1-\gamma}$ with $\gamma > 0$ the coefficient of relative risk aversion. They disagree on the dynamic of the expected production process, and are characterized by their own subjective beliefs, which give the subjective likelihood of occurrence of the different events in the economy. These subjective beliefs are indexed by δ, and, for a given $\delta \in \mathbb{R}$, the δ-subjective beliefs are defined by the subjective probability measure $\mathbb{Q}^{\delta}$, which is assumed to be equivalent to the true probability measure $\mathbb{P}$. I call Agent-δ the agent characterized by these beliefs. Concretely, all agents agree on the volatility of the production process σ but disagree on its drift.[3] Instead of considering that it equals μ, Agent-δ believes that the drift of the production process is given by $\mu+\delta$. Thus, δ represents her belief bias, and she is relatively optimistic (resp. pessimistic) compared to an agent with true beliefs if her bias is positive (resp. negative). Formally, Agent-δ believes that the expected production process is given

[3]This assumption is in line with the literature which shows that the expected return is harder to estimate than the variance (see, e.g., Williams, 1977, Merton, 1980).

by

$$dy_t = (\mu + \delta) \, y_t dt + \sigma y_t dW_{\delta,t},$$

where W_δ is a standard unidimensional $\left((\mathcal{F}_t), \mathbb{Q}^\delta\right)$-Brownian motion, such that $W_{\delta,t} = W_t - \frac{\delta t}{\sigma}$.

As there is a continuum of agents, I use a probability density function to describe their wealth share distribution, similarly to, e.g., Beddock and Jouini (2020). At time t and in state of the world W_t, I assume that the time-t wealth share distribution ν_{δ,t,W_t} of the continuum of investors is given by a normal probability density function with parameters $\bar{\delta}_{t,W_t}$ and $\bar{\omega}_{t,W_t}$, whose expression are endogenously determined later in the subsection. Thus, it is given by

$$\nu_{\delta,t,W_t} = \frac{1}{\sqrt{2\pi}\bar{\omega}_{t,W_t}} \exp\left(-\frac{\left(\delta - \bar{\delta}_{t,W_t}\right)^2}{2\bar{\omega}_{t,W_t}^2}\right).$$

At this time and in this state of the world, the investors consider a finite horizon economy of horizon T, and make their investment decisions. They forecast that the expected production process will deliver a payoff y_{t+T} at time $t+T$, and Agent-δ plans to consume $y_{\delta,t+T}$ at time $t+T$. The parameter T can therefore be associated to the agents' prevision horizon or, more generally, to the agents' investment horizon. Recall that in the equivalent overlapping heterogeneous generations model, this parameter T indicates the life span of each generation.

The time-t temporary equilibrium—defined by a positive density price p_{t+T} and a continuum of consumption plans $(y_{\delta,t+T})_{\delta\in\mathbb{R}}$, and given explicitly in Proposition 1—is obtained when each agent maximizes her expected utility from future consumption according to her time-t beliefs such that both her time-t static budget constraint and the market clearing condition are satisfied.

Let now study how this temporary equilibrium evolves as time goes by.

The main innovation of the model lies in the fact that, unlike similar existing models based on the well-known Arrow-Debreu framework (see, e.g., Jouini and Napp, 2007, Atmaz and Basak, 2018), I release the investor plan time consistency hypothesis. Roughly speaking, this implied assumption involves that, for all agents, today's actions coincide with yesterday's planned actions for the subsequent day, or, as stated by Lindahl (1939), it "implies that all plans prevailing at the starting point are based on expectations in conformity with reality, and that they will undergo no change with the lapse of time."[4] In such a setting, there is no need for the market to reopen after the starting date.

I thus release this assumption, and model an economy à la Lindahl (1939) with heterogeneous investors: I allow the agents to continuously adjust their consumption plans over time based on newly available information, and, for each agent, I do not require her continuum of plans to be time consistent. In the context of overlapping generations, this time inconsistency feature arises because successive generations have different information at their date of birth.

Formally, at time $t+dt$, the market reopens as new information W_{t+dt} comes in. Knowing this information, the agents update their actions. They still consider a finite horizon economy of horizon T. Thus, they forecast that the expected production process will deliver a payoff y_{t+T+dt} at time $t+T+dt$, and plan to consume at time $t+T+dt$. This leads to a time-$t+dt$ temporary equilibrium, defined by p_{t+T+dt} and $(y_{\delta,t+T+dt})_{\delta\in\mathbb{R}}$. As time goes by, the continuum of temporary equilibria therefore results in an infinite horizon economy where consumption plans are continuously revised and where effective consumption never occurs as it is continuously postponed. This is why I label the process y as the expected production process rather than the (effective) production process.[5]

The endowment of each agent being fixed, this setting implies an additional dynamic budget constraint. More precisely, seen from date t, and before trading and consumption reallocation, Agent-δ's expected consumption share should remain unchanged between $t+T$

[4]Erik Lindahl, *Studies in the theory of money and capital*, 1939: p.38

[5]In this model, I thus study planned consumption rather than effective consumption, and thereby extend to an infinite horizon framework this common feature of finite horizon models with consumption at a single (terminal) date (see, e.g., Atmaz and Basak, 2018). As mentioned in Section 2.2 and more formally explained in Appendix B, this issue can be circumvented by considering the alternative overlapping heterogeneous generations model that incorporates continuous effective consumption and yields the same implications.

and $t+T+dt$, which means that the expected value of her time-$t+T+dt$ consumption should equal the expected value of her adjusted time-$t+T$ consumption evaluated at date t and at the time-$t+T+dt$ price.[6] Formally, for each time t and each belief bias δ, this implies that

$$\mathbb{E}_t\left(p_{t+T+dt}y_{\delta,t+T+dt}\right)=\mathbb{E}_t\left(p_{t+T+dt}y_{\delta,t+T}\frac{y_{t+T+dt}}{y_{t+T}}\right). \tag{2.1}$$

Equation (2.1) simply states that the budget of each agent evolves according to the evolution of the total expected production. Solving this equation together with the time-t temporary equilibrium equations delivers explicit solutions for the time-t wealth share distribution parameters $\bar{\delta}_{t,W_t}$ and $\bar{\omega}_{t,W_t}$. They are given in Proposition 1, where I also report the equilibrium features, and give the characteristics of the representative agent, defined as the agent who, if endowed with the total wealth of the economy, would have a marginal utility equal to the equilibrium price.

Proposition 1 (Equilibrium and representative agent)**.**
In equilibrium, at time t and in state of the world W_t:

1. *The state price density and the investors' consumption plans are given by*

$$\begin{aligned}p_{t+T}&=y_{t+T}^{-\gamma}\left(\int\left(\lambda_{\delta,t,W_t}M_{\delta,t+T}\right)^{\frac{1}{\gamma}}d\delta\right)^{\gamma},\\ y_{\delta,t+T}&=y_{t+T}\left(\int\left(\lambda_{\delta,t,W_t}M_{\delta,t+T}\right)^{\frac{1}{\gamma}}d\delta\right)^{-1}\left(\lambda_{\delta,t,W_t}M_{\delta,t+T}\right)^{\frac{1}{\gamma}},\end{aligned}$$

where $M_{\delta,t+T}=exp\left(\frac{\delta}{\sigma}\left(W_{t+T}-W_t\right)-\frac{1}{2}\frac{\delta^2}{\sigma^2}T\right)$ is the time-t Radon-Nikodym derivative of the subjective probability measure $\mathbb{Q}^\delta$ with respect to $\mathbb{P}$, and

$$\lambda_{\delta,t,W_t}=\frac{1}{\sqrt{2\pi}^{\gamma}}\frac{\sqrt{\sigma^2+T\varphi\left(\bar{\omega}_{t,W_t}^2\right)}^{\gamma-1}}{\sqrt{\varphi\left(\bar{\omega}_{t,W_t}^2\right)\gamma}^{\gamma}\sqrt{\sigma^2}^{\gamma-1}}exp\left(-\frac{\left(\delta-\bar{\delta}_{t,W_t}+(1-\gamma)\,T\varphi\left(\bar{\omega}_{t,W_t}^2\right)\right)^2}{2\varphi\left(\bar{\omega}_{t,W_t}^2\right)}\right)$$

is the time-t Lagrange multiplier with

$$\varphi\left(x\right)=\frac{x}{2}-\frac{\gamma\sigma^2}{2T}+\sqrt{\left(\frac{x}{2}-\frac{\gamma\sigma^2}{2T}\right)^2+\frac{\sigma^2x}{T}}.$$

[6] Agent-δ's adjusted time-$t+T$ consumption is given by $y_{\delta,t+T}\frac{y_{t+T+dt}}{y_{t+T}}$, and takes into account the growth of the expected production process between $t+T$ and $t+T+dt$.

2. *The representative agent of the economy is the fictitious investor whose time-t Radon-Nikodym derivative of the subjective probability measure* $\mathbb{Q}^\delta$ *with respect to* $\mathbb{P}$ *is given by*

$$M_{RA,t+T} = \left(\int \left(\lambda_{\delta,t,W_t} M_{\delta,t+T} \right)^{\frac{1}{\gamma}} d\delta \right)^{\gamma}.$$

3. *The wealth share distribution of the continuum of investors is given by a normal probability density function with a standard deviation and a mean given respectively by*

$$\bar{\omega}_{t,W_t} = \bar{\omega}, \qquad \text{and} \qquad \bar{\delta}_{t,W_t} = \bar{\bar{\delta}} + \frac{\varphi\left(\bar{\omega}^2\right) W_t}{\sigma},$$

where $\bar{\bar{\delta}}$ *and* $\bar{\omega}$ *are given constants describing the time-*0 *wealth share distribution.*

From the third item, I observe that the distribution's standard deviation is constant and does not vanish as time goes by, which ensures that the agents differ persistently in their beliefs. This feature is in sharp contrast with Atmaz and Basak (2018) and other finite-horizon models with one terminal production date, where belief dispersion consistently decreases with time. Such a result is also of great empirical importance as surveys show that beliefs are mostly characterized by large and persistent individual heterogeneity (see, e.g., Meeuwis et al., 2019, Giglio et al., 2020, Das et al., 2020). Moreover, as the mean is state dependent, the average belief bias fluctuates around the initial value $\bar{\bar{\delta}}$. On average, the investors are more optimistic following good news, and more pessimistic following bad ones.

2.3 The stock price and its dynamics

I now derive the stock price and its dynamics in the presence of belief heterogeneity. To do so, I assume that a risky stock S is available for trading. The stock is in positive net supply of one unit and, at time t, is a claim to the payoff y_{t+T} expected to be paid at time $t+T$. Studying its properties leads to the following proposition.

Proposition 2 (Equilibrium stock price, mean return, and volatility)**.**
In equilibrium, at time t and in state of the world W_t*:*

1. The stock price is given by

$$S_t = \overline{S}_t exp\left(\bar{\delta}_{t,W_t}T - \frac{\varphi\left(\bar{\omega}^2\right)T^2}{2}\right),$$

2. The mean stock return is given by

$$\mu_{S_t} = \overline{\mu} + \sigma\left(\frac{\partial\bar{\delta}_{t,W_t}}{\partial W_t}T\right) + \frac{1}{2}\left(\frac{\partial\bar{\delta}_{t,W_t}}{\partial W_t}T\right)^2,$$

3. The stock volatility is given by

$$\sigma_{S_t} = \overline{\sigma} + \frac{\partial\bar{\delta}_{t,W_t}}{\partial W_t}T,$$

where $\overline{S}_t = y_t exp\left(\left(\mu - \gamma\sigma^2\right)T\right)$, $\overline{\mu} = \mu$, *and* $\overline{\sigma} = \sigma$ *are the equivalent quantities obtained in a similar standard economy without belief heterogeneity.*

The formulas share similarities with those in Atmaz and Basak (2018). As stated in the introduction, this is because my model builds on their work. However, the two models differ along one important dimension: unlike them, I design an economy where the effective consumption date is continuously postponed so that the time-t remaining time before consumption always equals T instead of $T - t$. As the time-t heterogeneity impact depends on the remaining time before consumption, I thus obtain a stationary model where the heterogeneity effects are not smoothed as time goes by. In particular, as the time-T stock price formula in Atmaz and Basak (2018) reduces to $S_T^{AB} = y_T$ (which is natural in their framework), the heterogeneity effects that they observe completely vanish when approaching date T, and their model is not stationary. Conversely, my time-T stock price is still fully impacted by the time-T investors heterogeneity, characterized by $\bar{\delta}_{T,W_T}$ and $\bar{\omega}$.

The results stationarity and persistence are of first interest as they more accurately reflect the heterogeneous market participants reality. These features also allow to test the model empirically over a long period. Before doing so in Section 3, I now discuss the properties of the market characteristics derived in Proposition 2. I first focus on the equilibrium mean stock return μ_{S_t}.

As stated before, one of the consequences of my methodology is that the time-t remaining time before expected effective consumption does not depend on t. Thus, the derivative of the stock price with respect to t differs markedly from the one obtained in Atmaz and Basak (2018), leading to different mean returns.

I find that a higher belief dispersion leads to a higher equilibrium mean stock return. More precisely, the belief dispersion has an impact on the sensitivity of the average belief bias to news: the higher the belief dispersion, the higher this sensitivity, and thus the higher the mean return. This is consistent with the recent work of Brandon and Wang (2020) who find that the average return on stocks with high sensitivity to earning belief shocks is 7.14% per year higher than that in stocks with low sensitivity. Conversely, Atmaz and Basak (2018) derive that the sign of the relation is state dependent and that higher dispersion leads to lower returns when the view is sufficiently optimistic. They further show that the relation between the mean return and the relative risk aversion coefficient depends on the level of optimism, while I derive an unambiguous negative relation. The intuition behind this result is simple: in an economy populated by heterogeneous agents, more risk averse agents speculate less aggressively, and thus earn lower returns. Lastly, I observe that the mean equilibrium stock return increases as the investment horizon T increases.

For the sake of completeness, I now briefly report the properties of the first and third items of Proposition 2, which are mostly similar to those in Atmaz and Basak (2018). I refer the reader to their paper for more detailed explanations of the underlying mechanisms behind these results.

Specifically, for a given time t, I find analogous impacts of the belief distribution parameters on the stock price. First, the stock price depends positively on the time-t average belief bias $\bar{\delta}_{t,W_t}$. Second, the sign of the belief dispersion impact is state dependent: the impact is positive for sufficiently good states of the world and negative for sufficiently bad ones.[7] I also similarly derive that the stock price is convex in the time-t expected pro-

[7] Two effects are at play. On the one hand, as the function φ is increasing in belief dispersion, there is a direct negative effect. On the other hand, there is an indirect effect of the belief dispersion trough

duction level y_t. Finally, unlike the standard case, the impact of the coefficient of relative risk aversion γ is not always negative, but can be positive for sufficiently bad states of the world.[8]

The third item considers the stock volatility σ_{S_t}. Several observations are in order. First, in a heterogeneous economy, it is higher than the production process volatility σ, in line with empirical observations (see, e.g., Ajinkya and Gift, 1985, Anderson et al., 2005, and Banerjee, 2011). An important difference with Atmaz and Basak (2018) is that, although both formulas have the same shape, this excess volatility effect does not decrease as time goes by. This is because the investors' heterogeneity remains persistent. Second, all else equal, the higher the belief dispersion, the higher the state-sensitivity of the average belief bias, and thus the higher the excess volatility. In fact, a higher fluctuations in the average belief bias translate to additional stock price fluctuations and therefore increases stock volatility. Additionally, the coefficient of relative risk aversion has a negative impact on σ_{S_t}. Finally, a higher investment horizon T leads to a higher stock volatility.

3 Empirical test of the model

As the model yields stationary formulas with persistent investor heterogeneity, the theoretical results can be empirically tested over long periods. Two additional observations are in order. First, as I consider an economy where there is only one risky stock available for trading, the model is better suited to test the market as a whole rather than to test individual stocks. Second, as observed in the previous section, the impact of the belief

the average belief bias. For bad states, this indirect effect is negative and reinforces the first effect. For good ones, this is the opposite and the overall effect can even be positive in case of sufficiently good states. Formally, as the average belief bias is state dependent, I derive that the stock price increases in belief dispersion when $\bar{\delta}_{t,W_t} > \bar{\delta} + \frac{\varphi\left(\bar{\omega}^2\right)T}{2}$.

[8] Formally, the stock price increases in risk aversion when the following condition holds

$$\bar{\delta}_{t,W_t} < \bar{\delta} + \frac{\varphi\left(\bar{\omega}^2\right)T}{2} - \sqrt{\left(\gamma\sigma^2 - T\bar{\omega}^2\right)^2 + 4T\sigma^2\bar{\omega}^2}.$$

bias on the model characteristics is indirect and depends on the belief dispersion. I thus focus on studying the empirical ability of belief dispersion to predict future market returns and volatility over time using variables sampled at the quarterly frequency from the first quarter of 1982 to the fourth quarter of 2019.

3.1 Market belief dispersion data

Let me first explain how the quarterly data on market belief dispersion, expressed in percentages, are constructed. Most importantly, building on the bottom-up measure of portfolio disagreement described in Yu (2011), I use analyst monthly forecasts as a proxy for investors' beliefs.

The data come from two databases: I use the analyst monthly forecasts of the EPS LTG of individual stocks from the IBES database,[9] and the CRSP database to obtain monthly market capitalizations. The IBES data are winsorized at the 1% and 99% levels to account for potential outliers or data errors. I also winsorize the prices at the 99% level. Furthermore, I exclude stocks with a price below five dollars at portfolio formation to avoid that extreme returns on penny stocks drive the results, and stocks for which less than two analysts provide EPS LTG forecasts during the month to focus on stocks that exhibit some forecast dispersion.

For each common stock i listed on the NYSE/Amex/Nasdaq in each month m, I obtain the standard deviation of the analyst forecasts—that I denote $\tilde{\omega}_{i,m}$—from the IBES Unadjusted Summary database.[10] Additionally, I obtain the market capitalization of each stock at the end of each month—that I denote $MKTCAP_{i,m}$—using the closing price and the

[9] I thus use earnings data to measure cashflows rather than dividends data. This choice is motivated by Da (2009) who argues that potential problems of working with dividends could arise because of the dividend payout policy of some firms. Campbell (2000) further highlights other empirical difficulties. On the theoretical side, using the accounting clean surplus identity, Vuolteenaho (1999) shows that if one looks at the infinite horizon, cash flow and earnings contain the same information. Thus, earnings are both theoretically equivalent and empirically better behaved than dividends.

[10] Similarly to Buraschi et al. (2014), I use unadjusted data to circumvent the problem of using stock-split adjusted data.

number of shares outstanding of the stock considered from CRSP. Similarly to Yu (2011), I thus construct a value-weighted measure of monthly market belief dispersion $\tilde{\omega}_m^{VW}$, defined as the cross-sectional (value-weighted) average of individual stock disagreement

$$\tilde{\omega}_m^{VW} = \frac{\sum_i \tilde{\omega}_{i,m} \times MKTCAP_{i,m}}{\sum_i MKTCAP_{i,m}}.$$

I then average the three monthly belief dispersions of a given quarter t to obtain quarterly market belief dispersion data $\bar{\omega}_t^{VW}$.

For the sake of robustness, I also construct a similar market belief dispersion variable, denoted by $\bar{\omega}_t^{EW}$, whose only difference with $\bar{\omega}_t^{VW}$ is to use equal-weighting rather than value-weighting when computing the monthly market dispersions. Figure 1 shows the time series of both variables.

Insert Figure 1 here.

In the remainder of the analysis, I therefore use $\bar{\omega}_t^{VW}$ and $\bar{\omega}_t^{EW}$ as the predictors, depending on whether the dependent variable of interest is constructed using value- or equal-weighting. I report their summary statistics over the full sample in Table 1. While both exhibit some persistence and are right-skewed, the value-weighted variable has a fatter tail than the equally-weighted one.

Insert Table 1 here.

3.2 Predicting market returns

Since the model abstracts from interest rate issues (as commonly done in models with no intermediate consumption), and thus does not allow to accurately define the risk premium, I focus on studying the ability of market belief dispersion to predict market (raw) returns. As implied by the second item of Proposition 2, a higher market belief dispersion should lead to higher market returns. This question is of first importance when dealing, e.g., with portfolio management or capital budgeting.

I first compute the returns of the overall market, which is defined each month as the portfolio containing all individual stocks for which at least two monthly EPS LTG forecasts are provided in the IBES database. Hence, the assets that constitute the market are those used to construct the market belief dispersion variables $\bar{\omega}^{VW}$ and $\bar{\omega}^{EW}$, which allows to more precisely capture the link between the market characteristics and the investor beliefs.

Insert Figure 2 here.

Similarly to the belief dispersion variables, I define RET^{VW} and RET^{EW}, the annualized market holding returns from one quarter to the next one, obtained from a value-weighted and an equally-weighted market portfolio respectively. I plot the variables in Figure 2, and report their summary statistics in Table 2. Note that, as underlined by Albuquerque (2012), the returns of the aggregate market are negatively skewed.

Insert Table 2 here.

I then test the model implied positive relation: a higher market belief dispersion in a given quarter should result in higher market returns in the subsequent quarter. To do so, I run the following standard OLS regression

$$RET_t^i = \gamma^i + \theta^i \bar{\omega}_{t-1}^i + \xi_t^i, \tag{3.1}$$

where t refers to quarter t, and $i = \{\text{VW, EW}\}$. Inference is based on autocorrelation- and heteroskedasticity-robust standard errors (Newey and West, 1987), and all variables are standardized prior to estimation.

Exploiting the whole sample of data, the implication of the model seems to be verified: using value-weighted variables yields a coefficient of 0.09 associated to a t-statistics of 1.39, while the equally-weighted specification leads to $\hat{\theta}^{EW} = 0.13$ with a t-statistics of 1.98.

In order to gain deeper insights of this predicted positive relation, I further run rolling window regressions of Equation (3.1) using subsamples of 15 years, i.e. with 60 quarterly observations. I thus look at the evolution of the predictability of market returns by market belief dispersion over time. More precisely, the framework leads to a total of 92 regressions for each specification,[11] and I therefore obtain time series of the estimated values of θ^{VW} and θ^{EW}, that are reported in Figure 3. In each panel, the horizontal axis shows the end date of the subsamples, while the vertical axis gives the estimated value. Thicker rounds (resp. crosses) indicate statistically significant positive (resp. negative) values at the 10% level.

Insert Figure 3 here.

The figure shows that all subsamples ending between 2011 and 2019 yield positive estimated values for both types of weighting, while subsamples ending during the previous decade lead to negative ones. This is especially true for the value-weighted specification: 12 regressions result in significantly positive $\hat{\theta}^{VW}$ in the last decade, and 11 significantly negative coefficients are obtained from regressions ending between 2001 and 2010. This might explain why the coefficient is not significantly different from zero using all data.

Finally, Table 3 reports the number of positive estimated coefficients for both specifications, along with the number of significantly positive (second column) and negative (third column) ones. Panel A presents the results for rolling windows of 60 quarterly observations, and Panels B and C use different window lengths of 40 and 80 quarterly observations respectively. Apart from the value-weighted specification of Panel B, the results further confirm that the model implication seems to be verified empirically.

Insert Table 3 here.

3.3 Predicting market volatility

The question of volatility forecasting is also of great empirical importance. For instance, from a risk-management perspective, accurate predictions of future volatility help to compute value-at-risk over long horizons. I thus test if market belief dispersion predicts market volatility at a quarterly horizon, and study how this predictive relation evolves over time. More precisely, following the implication of the third item of Proposition 2, I investigate if

[11]The first subsample starts in the second quarter of 1982 and ends in the first quarter of 1997.

a higher market belief dispersion in a given quarter leads to a higher market volatility in the subsequent quarter.

As commonly done in the literature (see, e.g., French et al., 1987, Schwert, 1989), I exploit daily stock returns to obtain my quarterly market volatility data, and, as before, define two variables using either value- or equal-weighting. Let focus on the construction of the value-weighted variable. For each month of a given quarter, I define a value-weighted portfolio of all individual stocks for which at least two analyst monthly forecasts of the EPS LTG are available on the IBES database. Again, this methodology ensures that the belief dispersion variables capture the investors' beliefs concerning the assets used in the construction of the dependent variables, and strengthens the link between them. Using daily stock return data from CRSP, I then compute the squared daily returns of these three monthly value-weighted market portfolios, and, summing them over the quarter (after subtracting the average daily market return in the quarter), I obtain the quarterly value-weighted market variance that I easily convert into the annualized quarterly value-weighted market volatility, VOL^{VW}. Computing its descriptive statistics shows that the variable is highly positively skewed and leptokurtic.[12] Similarly to Paye (2012), I thus define the annualized quarterly market log volatility variable $LVOL^{VW}$ as the natural logarithm of VOL^{VW}, whose distribution is approximately Gaussian (Andersen et al., 2001). Because the empirical analysis relies on linear models estimated by OLS, this latter property is of first importance, and I therefore use $LVOL^{VW}$ as the independent variable in the subsequent empirical analysis.

Insert Figure 4 here.

The annualized quarterly market equally-weighted log volatility variable $LVOL^{EW}$ is constructed analogously using an equally-weighted portfolio each month, and Figure 4 shows the time series of both variables.

Insert Table 4 here.

[12]Over the entire sample, the skewness of VOL^{VW} is 2.89, and its kurtosis equals 15.48.

I also report their summary statistics in Table 4. Unsurprisingly, both series exhibit a high degree of persistence, which is a known characteristic of volatility processes. Unreported graphs of their sample autocorrelation and partial autocorrelation functions further suggest that they are consistent with AR processes, and the Bayesian information criterion advocates the choice of two lags in the dependent variables. As in Paye (2012), I thus consider the following regression for log volatility

$$LVOL_t^i = \alpha_0^i + \alpha_1^i LVOL_{t-1}^i + \alpha_2^i LVOL_{t-2}^i + \beta^i \bar{\omega}_{t-1}^i + \epsilon_t^i, \tag{3.2}$$

where t refers to quarter t, and $i = \{\text{VW, EW}\}$.

The main interest of this standard OLS regression is to test the hypothesis $H_0 : \beta = 0$ against the alternative $H_1 : \beta \neq 0$: rejecting the null indeed implies that, when controlling for past (log) volatility, belief dispersion helps to predict future market (log) volatility. Similarly to the previous part, inference is based on autocorrelation- and heteroskedasticity-robust standard errors, and all variables are standardized prior to estimation.

I first run the regression defined in (3.2) using all data available for value- and equally-weighted variables. The results show that β is positive for both specifications although it is not statistically different from zero.[13] Thus, over the full sample, the theoretical implication of the model presented in Section 2 regarding stock volatility does not seem to be strongly validated by the data.

However, to further investigate this relation and to study how the ability of belief dispersion to predict market (log) volatility evolves over time, I run rolling window regressions of Equation (3.2) using subsamples of 15 years.[14] This leads to a total of 91 regressions for each specification.[15] I therefore obtain time series of the estimated values of β^{VW} and β^{EW}, and report them in Figure 5, which is constructed in the same fashion as Figure 3.

Insert Figure 5 here.

[13]More precisely, I find $\beta^{VW} = 0.04$ with a t-stat of 0.78, and $\beta^{EW} = 0.01$ with a t-stat of 0.15.

[14]Conrad and Glas (2018) provide a similar analysis to test if macroeconomic variables predict volatility in the cross-section of industry portfolios.

[15]Because I need to have the market (log) volatility data of the two previous quarters, my first subsample starts in the third quarter of 1982 and ends in the second quarter of 1997.

As reported in Panel A of Table 5, 83 (resp. 52) of the 91 estimated β's are positive when using value-weighted (resp. equally-weighted) variables, among which 17 (resp. 22) are statistically significant at the 10% level which corresponds to roughly a fifth (resp. a quarter) of the cases considered. Moreover, none of the negative coefficients is significantly different from zero. The evidence therefore points towards the approval of the model-implied positive effect of belief dispersion on the market volatility. Note, however, that the predictive ability of the market belief dispersion varies over time: most of the significantly positive β's are obtained for subsamples ending between 2000 and 2010, and the negative ones are concentrated in the most recent subsamples.

Insert Table 5 here.

As robustness checks, Panels B to E of Table 5 further show results of alternative specifications of the rolling window regressions. Panel B (resp. Panel C) considers the same specification adding (resp. retrieving) one lagged value of (log) volatility to Equation (3.2). Panels D and E consider rolling window regressions using alternative subsamples sizes (10 and 20 years respectively). Overall, they confirm the results found in the main specification of the analysis.

4 Conclusion

Building on the work of Atmaz and Basak (2018), and allowing for consumption plan revisions in the spirit of Lindahl (1939), I model an infinite horizon economy populated by a large number of investors who differ in their beliefs. Specifically, I use a Gaussian distribution with a state-dependent mean and a constant standard deviation to describe the investors' wealth shares. After solving the model—which is fully equivalent to an overlapping heterogeneous generations model that can incorporates effective consumption—I study the implications of belief heterogeneity on various quantities of interest, namely the stock price, its mean return, and its volatility. In particular, I derive that both the stock mean return and volatility monotonically increase with belief dispersion. While the derived formulas share similarities with those in Atmaz and Basak (2018), a major difference is that the theoretical framework that I use leads to stationary results with non-vanishing heterogeneity. The interest in such a modeling is twofold. First, the heterogeneity persistence is empirically observed, and consequently the model results are more in line with reality. Second, it allows me to test the model empirically over long periods, and to assess its relevance. Using analyst forecasts from the IBES database, I thus show evidence that the documented positive relations between the stock equilibrium mean return and volatility and the belief dispersion are verified in the data when considering the market as a whole.

Note that the model only considers a single stock in the economy. It would thus be interesting to extend it to the case of a multi-stocks economy to derive testable cross-sectional relations. I leave this for future research.

A Proofs of the propositions

Proof of Proposition 1

1. 2. Formally, at time t and in state of the world W_t, Agent-δ's maximization program is given by

$$\max_{y_{\delta,t+T}} \mathbb{E}_t\left(M_{\delta,t+T}u\left(y_{\delta,t+T}\right)\right) \quad \text{such that} \quad \left(\nu_{\delta,t,W_t}y_{t+T}-y_{\delta,t+T}\right)p_{t+T}\geq 0, \tag{A.1}$$

where $M_{\delta,t+T}$ is the time-t Radon-Nikodym derivative of the subjective probability measure $\mathbb{Q}^{\delta}$ with respect to $\mathbb{P}$, such that

$$M_{\delta,t+T}=\exp\left(\frac{\delta}{\sigma}\left(W_{t+T}-W_t\right)-\frac{1}{2}\frac{\delta^2}{\sigma^2}T\right).$$

Moreover, the market clearing condition is explicitly given by

$$y_{t+T}=\int y_{\delta,t+T}d\delta. \tag{A.2}$$

Solving Program (A.1) such that Equation (A.2) is satisfied leads to

$$\begin{cases} p_{t+T} &= y_{t+T}^{-\gamma}\left(\int\left(\lambda_{\delta,t,W_t}M_{\delta,t+T}\right)^{\frac{1}{\gamma}}d\delta\right)^{\gamma}, \\ y_{\delta,t+T} &= y_{t+T}\left(\int\left(\lambda_{\delta,t,W_t}M_{\delta,t+T}\right)^{\frac{1}{\gamma}}d\delta\right)^{-1}\left(\lambda_{\delta,t,W_t}M_{\delta,t+T}\right)^{\frac{1}{\gamma}}, \end{cases}$$

where λ_{δ,t,W_t} is the time-t Lagrange multipliers of the form

$$\lambda_{\delta,t,W_t}=K_{t,W_t}\exp\left(-\frac{1}{2a_{t,W_t}^2}\delta^2-b_{t,W_t}\delta\right). \tag{A.3}$$

To further identify a_{t,W_t} and b_{t,W_t}, notice that, by definition, the following equation must hold

$$\nu_{\delta,t,W_t}=\frac{\mathbb{E}_t\left(p_{t+T}y_{\delta,t+T}\right)}{\mathbb{E}_t\left(p_{t+T}y_{t+T}\right)}.$$

By identification, this leads to

$$\begin{cases} a_{t,W_t}^2 &= \frac{\bar{\omega}_{t,W_t}^2}{2}-\frac{\gamma\sigma^2}{2T}+\sqrt{\left(\frac{\bar{\omega}_{t,W_t}^2}{2}-\frac{\gamma\sigma^2}{2T}\right)^2+\frac{\sigma^2\bar{\omega}_{t,W_t}^2}{T}}, \\ b_{t,W_t} &= \left(1-\gamma\right)T-\frac{\bar{\delta}_{t,W_t}}{a_{t,W_t}^2}. \end{cases}$$

Finally, in order to identify K_{t,W_t}, I define the representative investor of the economy. By definition, this agent would have a marginal utility equal to the equilibrium price if she were endowed with the total wealth of the economy, such that $p_{t+T} = y_{t+T}^{-\gamma} M_{AR,t+T}$. By identification, her time-t subjective beliefs' Radon-Nikodym derivative with respect to the true beliefs $\mathbb{P}$ is thus given by

$$M_{AR,t+T} = \left(\int \left(\lambda_{\delta,t,W_t} M_{\delta,t+T} \right)^{\frac{1}{\gamma}} d\delta \right)^{\gamma}.$$

I need to ensure that $\mathbb{E}_t \left(M_{AR,t+T} \right) = 1$ for the representative agent beliefs to be well-defined. Thus, after some easy computations, K_{t,W_t} is given by

$$K_{t,W_t} = \frac{1}{\sqrt{2\pi}^{\gamma}} \frac{\sqrt{\sigma^2 + T a_{t,W_t}^2}^{\gamma-1}}{\sqrt{a_{t,W_t}^2 \gamma}^{\gamma} \sqrt{\sigma^2}^{\gamma-1}} \exp\left(-\frac{a_{t,W_t}^2 b_{t,W_t}^2}{2} \right).$$

Defining the function φ as given in Proposition 1 and plugging the expressions of a_{t,W_t}, b_{t,W_t}, and K_{t,W_t} into Equation (A.3) yield the results.

3. Let define $\mu_{p_{t+T}}$ and $\sigma_{p_{t+T}}$ such that

$$dp_{t+T} = \mu_{p_{t+T}} p_{t+T} dt + \sigma_{p_{t+T}} p_{t+T} dW_t.$$

$\mu_{y_{\delta,t+T}}$ and $\sigma_{y_{\delta,t+T}}$ are defined similarly. Note that these four quantities depend on $\bar{\delta}_{t,W_t}$ and $\bar{\omega}_{t,W_t}$.

Using Ito's lemma, Equation (2.1) leads to

$$\mathbb{E}_t \left(\mu_{y_{\delta,t+T}} + \sigma_{p_{t+T}} \sigma_{y_{\delta,t+T}} - \mu - \sigma \sigma_{p_{t+T}} \right) = 0. \tag{A.4}$$

Direct computations allow to rewrite the left hand-side of Equation (A.4) as a polynomial function of δ of degree two. As Equation (A.4) must be verified for all agents, each coefficient of the polynomial form must equal zero. By identification, this leads to the expressions of $\bar{\delta}_{t,W_t}$ and $\bar{\omega}_{t,W_t}$ derived in the proposition.

Proof of Proposition 2

1. By no arbitrage, the time-t stock price is given by

$$S_t = \frac{\mathbb{E}_t\left(p_{t+T}y_{t+T}\right)}{\mathbb{E}_t\left(p_{t+T}\right)}.$$

   Computing the numerator and the denominator and rearranging the terms lead to the formula in Proposition 2.

   To determine the benchmark economy stock price $\overline{S}_t$, I set $\bar{\delta}$ and $\bar{\omega}$ to zero, and substitute them into the stock price formula.

2. 3. Applying Ito's lemma to the time-t stock price formula yields the results.

   Similarly to the first item, $\overline{\mu}$ and $\overline{\sigma}$ are defined by setting $\bar{\delta}$ and $\bar{\omega}$ to zero into the formulas.

B Equivalence of the two models

Let denote by $\tilde{y}_{\delta,t+T}$ the endowment of Agent-δ of Generation-t in the overlapping heterogeneous generations model. She consumes $c_{\delta,t+T}$ and bequests $b_{\delta,t+T}$, such that $\tilde{y}_{\delta,t+T} = c_{\delta,t+T} + b_{\delta,t+T}$. Her utility is thus given by

$$\mathbb{E}_t\left(M_{\delta,t+T}\left(au\left(c_{\delta,t+T}\right) + u\left(b_{\delta,t+T}\right)\right)\right), \tag{B.1}$$

where a is a given non-negative coefficient, common to all agents and all generations, that represents the agents' degree of selfishness.

Maximizing Expression (B.1), under the condition that the sum of consumption and bequest equals the endowment, leads to

$$c_{\delta,t+T} = \frac{1}{1+a^{-\frac{1}{\gamma}}}\tilde{y}_{\delta,t+T}, \qquad \text{and} \qquad b_{\delta,t+T} = \frac{a^{-\frac{1}{\gamma}}}{1+a^{-\frac{1}{\gamma}}}\tilde{y}_{\delta,t+T}.$$

Maximizing Expression (B.1) is therefore equivalent to finding the endowment $\tilde{y}_{\delta,t+T}$ that maximizes $\mathbb{E}_t\left(M_{\delta,t+T}u\left(\tilde{y}_{\delta,t+T}\right)\right)$, under the usual market clearing condition and budget constraint, before dividing it into consumption and bequest according to the above mentioned proportions.

Let further define $\frac{1}{1+a^{-\frac{1}{\gamma}}} = \alpha dt$. Because the overall endowment $\tilde{y}$ is partly consumed at each period, the overall endowment process is given by

$$d\tilde{y}_t = (\tilde{\mu} - \alpha)\,\tilde{y}_t dt + \tilde{\sigma}\tilde{y}_t dW_t.$$

Thus, the maximization program of this heterogeneous overlapping generations model is equivalent to Program (A.1). The two programs yield identical results by setting $\tilde{\mu} = \mu + \alpha$ and $\tilde{\sigma} = \sigma$, and the two models therefore lead to the same implications.

Note that setting $a = 0$ corresponds to the case without effective consumption, or, stated differently, to the case where each generation bequests all its endowment to the next one. This specific case leads to $\alpha = 0$, and thus $\tilde{\mu} = \mu$.

References

Bipin B. Ajinkya and Michael J. Gift. Dispersion of financial analysts' earnings forecasts and the (option model) implied standard deviations of stock returns. The Journal of Finance, 40(5):1353–1365, 1985.

Rui Albuquerque. Skewness in stock returns: Reconciling the evidence on firm versus aggregate returns. Review of Financial Studies, 25(5):1630–1673, 2012.

Torben G. Andersen, Tim Bollerslev, Francis X. Diebold, and Heiko Ebens. The distribution of realized stock return volatility. Journal of Financial Economics, 61(1):43–76, 2001.

Evan W. Anderson, Eric Ghysels, and Jennifer L. Juergens. Do heterogeneous beliefs matter for asset pricing? Review of Financial Studies, 18(3):875–924, 2005.

Daniel Andrei, Bruce Carlin, and Michael Hasler. Asset pricing with disagreement and uncertainty about the length of business cycles. Management Science, 65(6):2900–2923, 2019.

Adem Atmaz and Suleyman Basak. Belief dispersion in the stock market. The Journal of Finance, 73(3):1225–1279, 2018.

Doron Avramov, Tarun Chordia, Gergana Jostova, and Alexander Philipov. Dispersion in analysts' earnings forecasts and credit rating. Journal of Financial Economics, 91(1): 83–101, 2009.

Snehal Banerjee. Learning from prices and the dispersion in beliefs. Review of Financial Studies, 24(9):3025–3068, 2011.

Snehal Banerjee and Ilan Kremer. Disagreement and learning: Dynamic patterns of trade. The Journal of Finance, 65(4):1269–1302, 2010.

Sudipta Basu. The conservatism principle and the asymmetric timeliness of earnings. Journal of Accounting and Economics, 24(1):3–37, 1997.

Arthur Beddock and Elyès Jouini. Live fast, die young: equilibrium and survival in large economies. Economic Theory, 2020.

Rajna Gibson Brandon and Songtao Wang. Earnings belief risk and the cross-section of stock returns. Review of Finance, 24(5):1107–1158, 2020.

Andrea Buraschi and Alexei Jiltsov. Model uncertainty and option markets with heterogeneous beliefs. The Journal of Finance, 61(6):2841–2897, 2006.

Andrea Buraschi, Fabio Trojani, and Andrea Vedolin. Economic uncertainty, disagreement, and credit markets. Management Science, 60(5):1281–1296, 2014.

John Y. Campbell. Asset pricing at the millennium. The Journal of Finance, 55(4): 1515–1567, 2000.

Joseph Chen, Harrison Hong, and Jeremy C. Stein. Breadth of ownership and stock returns. Journal of Financial Economics, 66(2-3):171–205, 2002.

Christian Conrad and Alexander Glas. 'Déjà vol' revisited: Survey forecasts of macroeconomic variables predict volatility in the cross-section of industry portfolios. SSRN Working paper, 2018. URL https://ssrn.com/abstract=3186567.

Jennifer Conrad, Bradford Cornell, and Wayne R. Landsman. When is bad news really bad news? The Journal of Finance, 57(6):2507–2532, 2002.

Jaksa Cvitanic and Semyon Malamud. Price impact and portfolio impact. Journal of Financial Economics, 100(1):201–225, 2011.

Zhi Da. Cash flow, consumption risk, and the cross-section of stock returns. The Journal of Finance, 64(2):923–956, 2009.

Sreyoshi Das, Camelia M Kuhnen, and Stefan Nagel. Socioeconomic status and macroeconomic expectations. Review of Financial Studies, 33(1):395–432, 2020.

Karl B. Diether, Christopher J. Malloy, and Anna Scherbina. Differences of opinion and the cross section of stock returns. The Journal of Finance, 57(5):2113–2141, 2002.

John A. Doukas, Chansog (Francis) Kim, and Christos Pantzalis. Divergence of opinion and equity returns. Journal of Financial and Quantitative Analysis, 41(3):573–606, 2006.

Kenneth R. French, G. William Schwert, and Robert F. Stambaugh. Expected stock returns and volatility. Journal of Financial Economics, 19(1):3–29, 1987.

Stefano Giglio, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus. Five facts about beliefs and portfolios. SSRN Working paper, 2020. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3336400.

Jean Michel Grandmont. Temporary general equilibrium theory. Econometrica, 45(3):535, 1977.

Jean-Michel Grandmont. Temporary Equilibrium, chapter Temporary Equilibrium, pages 6579–6583. Palgrave Macmillan UK, London, 2008.

Elyès Jouini and Clotilde Napp. Consensus consumer and intertemporal asset pricing with heterogeneous beliefs. Review of Economic Studies, 74(4):1149–1174, 2007.

Elyès Jouini and Clotilde Napp. Unbiased disagreement in financial markets, waves of pessimism and the risk-return trade-off. Review of Finance, 15(3):575–601, 2011.

Mordecai Kurz and Maurizio Motolese. Diverse beliefs and time variability of risk premia. Economic Theory, 47(2-3):293–335, 2011.

Erik Lindahl. Studies in the theory of money and capital. London, 1939.

Maarten Meeuwis, Jonathan Parker, Antoinette Schoar, and Duncan Simester. Belief disagreement and portfolio choice. 2019. doi: 10.3386/w25108.

Robert C. Merton. On estimating the expected return on the market. Journal of Financial Economics, 8(4):323–361, 1980.

Edward M. Miller. Risk, uncertainty, and divergence of opinion. The Journal of Finance, 32(4):1151–1168, 1977.

Stefan Nagel. Short sales, institutional investors and the cross-section of stock returns. Journal of Financial Economics, 78(2):277–309, 2005.

Whitney K. Newey and Kenneth D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica, 55(3):703–708, 1987.

Bradley S. Paye. 'Déjà vol': Predictive regressions for aggregate stock market volatility using macroeconomic variables. Journal of Financial Economics, 106(3):527–546, 2012.

José A. Scheinkman and Wei Xiong. Overconfidence and speculative bubbles. Journal of Political Economy, 111(6):1183–1220, 2003.

G. William Schwert. Why does stock market volatility change over time? The Journal of Finance, 44(5):1115–1153, 1989.

Catherine T. Shalen. Volume, volatility, and the dispersion of beliefs. Review of Financial Studies, 6(2):405–434, 1993.

Pietro Veronesi. Stock market overreactions to bad news in good times: A rational expectations equilibrium model. Review of Financial Studies, 12(5):975–1007, 1999.

Tuomo Vuolteenaho. Understanding the aggregate book-to-market ratio. 1999.

Joseph T. Williams. Capital asset prices with heterogeneous beliefs. Journal of Financial Economics, 5(2):219–239, 1977.

Jianguo Xu. Price convexity and skewness. The Journal of Finance, 62(5):2521–2552, 2007.

Jialin Yu. Disagreement and return predictability of stock portfolios. Journal of Financial Economics, 99(1):162–183, 2011.

Figure 1: Time series of the market belief dispersion variables

Panel A: Value-weighted variable

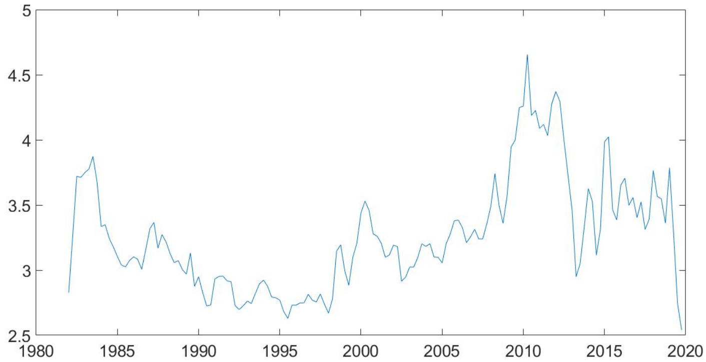

Panel B: Equally-weighted variable

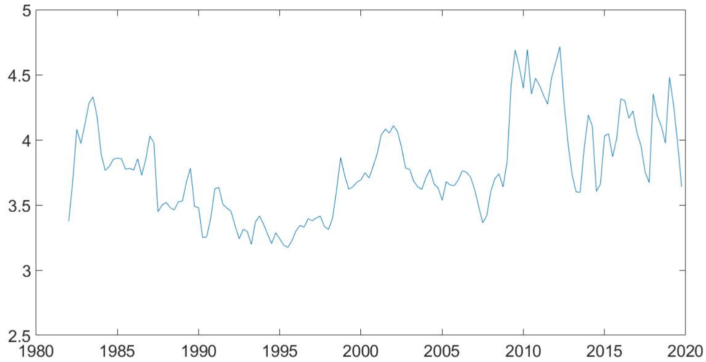

The figure presents the time evolution of the quarterly market belief dispersion variables $\bar{\omega}^{VW}$ *(Panel A) and* $\bar{\omega}^{EW}$ *(Panel B) expressed in percentages. The sample goes from the first quarter of 1982 to the fourth quarter of 2019.*

Figure 2: Time series of the annualized quarterly market holding return variables

Panel A: Value-weighted variable

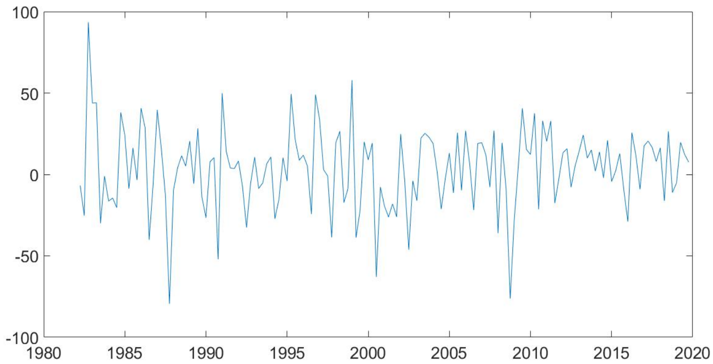

Panel B: Equally-weighted variable

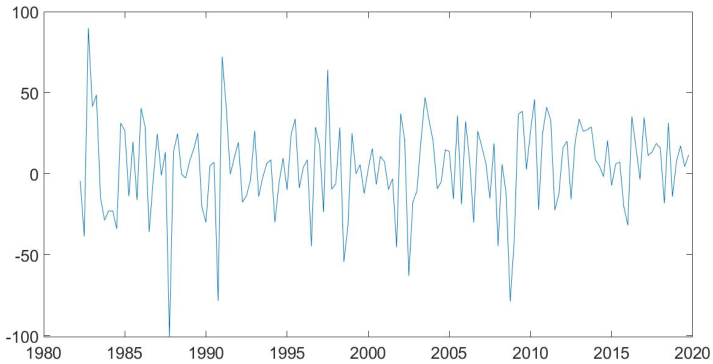

The figure presents the time evolution of the annualized quarterly market return variables RET^{VW} *(Panel A) and* RET^{EW} *(Panel B) expressed in percentages. The sample goes from the second quarter of 1982 to the fourth quarter of 2019.*

Figure 3: Time series of rolling window estimates of θ

Panel A: Value-weighted specification

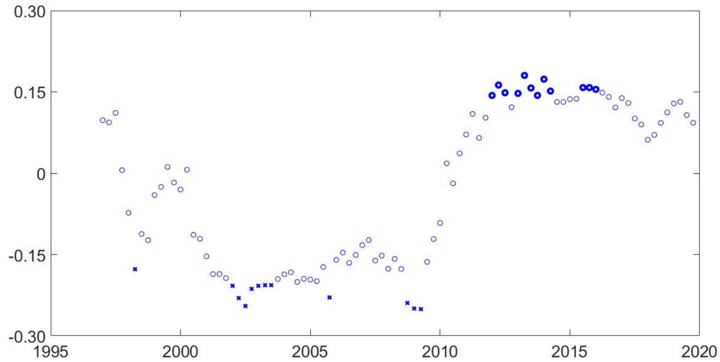

Panel B: Equally-weighted specification

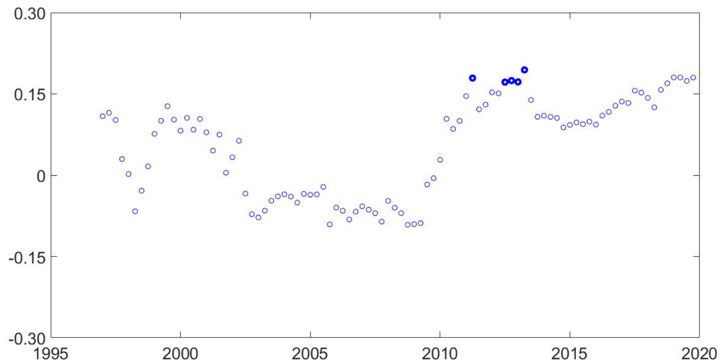

The figure presents the estimated values of θ^{VW} (Panel A) and θ^{EW} (Panel B) obtained from rolling window regressions with 60 quarterly observations. The horizontal axis shows the end date of the subsamples (the first one ends in the first quarter of 1997). Thicker rounds (resp. crosses) indicate statistically significant positive (resp. negative) values at the 10% level.

Figure 4: Time series of the annualized quarterly market log volatility variables

Panel A: Value-weighted variable

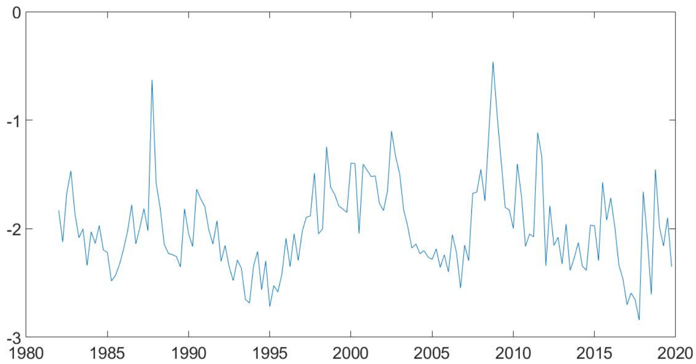

Panel B: Equally-weighted variable

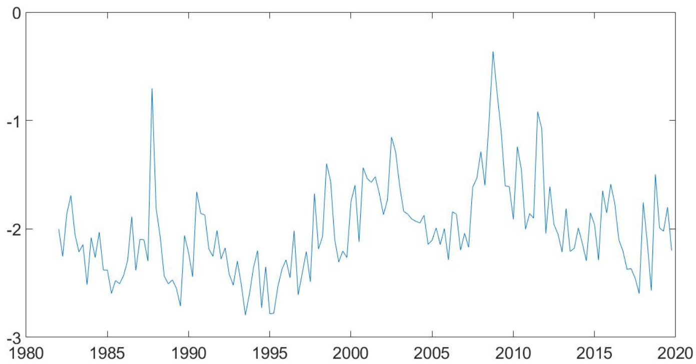

The figure presents the time evolution of the annualized quarterly market log volatility variables $LVOL^{VW}$ (Panel A) and $LVOL^{EW}$ (Panel B). The sample goes from the first quarter of 1982 to the fourth quarter of 2019.

Figure 5: Time series of rolling window estimates of β

Panel A: Value-weighted specification

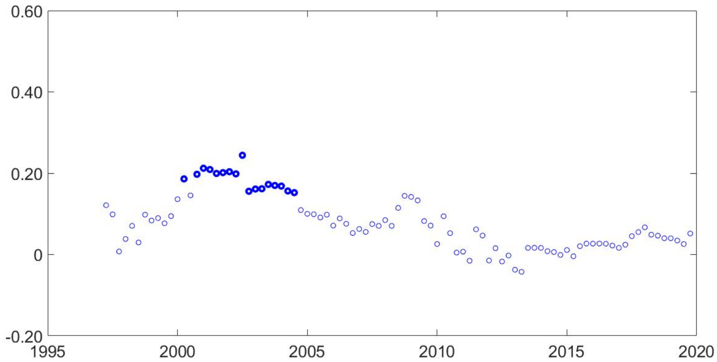

Panel B: Equally-weighted specification

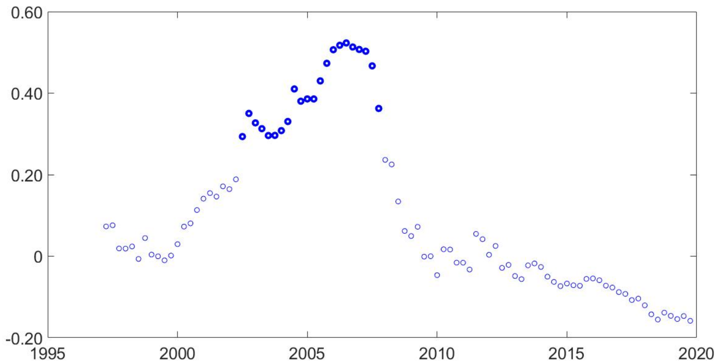

The figure presents the estimated values of β^{VW} (Panel A) and β^{EW} (Panel B) obtained from rolling window regressions with 60 quarterly observations. The horizontal axis shows the end date of the subsamples (the first one ends in the second quarter of 1997). Thicker rounds indicate statistically significant positive values at the 10% level.

Table 1: Descriptive statistics of the market belief dispersion variables

	Mean	St. Dev.	Skewness	Kurtosis	ρ_1	ρ_2
$\bar{\omega}^{VW}$	3.32	0.51	0.91	3.30	0.85	0.75
$\bar{\omega}^{EW}$	3.78	0.37	0.53	2.63	0.88	0.74

The table contains descriptive statistics for the two belief dispersion variables $\bar{\omega}^{VW}$ and $\bar{\omega}^{EW}$, expressed in percentages, considered in the paper. The mean, standard deviation, skewness, and kurtosis are reported for each variable, as well as the first- and second-order sample autocorrelations (ρ_1 and ρ_2). The variables are sampled at a quarterly frequency from the first quarter of 1982 to the fourth quarter of 2019.

Table 2: Descriptive statistics of the annualized quarterly market holding return variables

	Mean	St. Dev.	Skewness	Kurtosis
RET^{VW}	3.43	0.25	-0.25	4.47
RET^{EW}	4.15	0.28	-0.53	4.49

The table contains descriptive statistics for the two annualized quarterly market return variables RET^{VW} and RET^{EW}, expressed in percentages, considered in the paper. The mean, standard deviation, skewness, and kurtosis are reported for each variable. The variables are sampled at a quarterly frequency from the second quarter of 1982 to the fourth quarter of 2019.

Table 3: Statistics of rolling window estimates of θ^i

i	$\#\theta > 0$	$\#\theta^* > 0$	$\#\theta^* < 0$
	Panel A: 60 quarterly data (92 samples)		
VW	44	12	12
EW	60	5	0
	Panel B: 40 quarterly data (112 samples)		
VW	55	1	16
EW	75	7	0
	Panel C: 80 quarterly data (72 samples)		
VW	40	9	5
EW	59	2	0

The table contains some statistics on the estimated values of θ^i (i = VW,EW) obtained from rolling window regressions given by Equation (3.1). *Panel A (resp. Panel B, Panel C) uses rolling windows of 60 (resp. 40, 80) quarterly data. The first column reports the number of positive coefficients obtained from these regressions. The second (resp. third) column report the number of significantly positive (resp. negative) coefficients at the* 10% *level obtained from these regressions. The overall sample goes from the second quarter of 1982 to the fourth quarter of 2019.*

Table 4: Descriptive statistics of the annualized quarterly market log volatility variables

	Mean	St. Dev.	Skewness	Kurtosis	ρ_1	ρ_2	ρ_3
$LVOL^{VW}$	-2.00	0.41	0.81	4.13	0.60	0.18	0.14
$LVOL^{EW}$	-2.01	0.43	0.92	4.29	0.61	0.21	0.16

The table contains descriptive statistics for the two annualized quarterly market log volatility variables $LVOL^{VW}$ and $LVOL^{EW}$ considered in the paper. The mean, standard deviation, skewness, and kurtosis are reported for each variable, as well as the first-, second-, and third-order sample autocorrelations (ρ_1, ρ_2, and ρ_3). The variables are sampled at a quarterly frequency from the first quarter of 1982 to the fourth quarter of 2019.

Table 5: Statistics of rolling window estimates of β^i

i	$\#\beta > 0$	$\#\beta^* > 0$	$\#\beta^* < 0$
Panel A: 60 quarterly data and 2 lags (91 samples)			
VW	83	17	0
EW	52	22	0
Panel B: 60 quarterly data and 1 lag (92 samples)			
VW	86	19	0
EW	59	29	0
Panel C: 60 quarterly data and 3 lags (90 samples)			
VW	76	6	0
EW	42	15	6
Panel D: 40 quarterly data and 2 lags (111 samples)			
VW	92	22	0
EW	59	27	4
Panel E: 80 quarterly data and 2 lags (71 samples)			
VW	53	12	0
EW	52	12	0

The table contains some statistics on the estimated values of β^i (i = VW,EW) obtained from rolling window regressions for several settings. Panel A refers to the specification defined by Equation (3.2) *and uses 60 quarterly data. Panel B (resp. Panel C) adds (resp. retrieves) one lag to this specification. Panel D and E show results when the rolling window regressions use 40 quarterly data and 80 quarterly data respectively. The first column reports the number of positive coefficients obtained from these regressions. The second (resp. third) column report the number of significantly positive (resp. negative) coefficients at the* 10% *level obtained from these regressions. The overall sample goes from the first quarter of 1982 to the fourth quarter of 2019.*