Mathematical Social Sciences 125 (2023) 27–41

Contents lists available at ScienceDirect

Mathematical Social Sciences

journal homepage: www.elsevier.com/locate/mss

Unrealized arbitrage opportunities in naive equilibria with non-Bayesian belief processes ☆

Alexander Zimper

Department of Economics, University of Pretoria, South Africa

ARTICLE INFO

Article history:
Received 14 July 2022
Received in revised form 2 July 2023
Accepted 3 July 2023
Available online 11 July 2023

Keywords:
Non-Bayesian updating
Dynamic inconsistency
Bounded rationality
Representative agent
Dynamic arbitrage opportunities

ABSTRACT

A non-Bayesian decision maker forms posterior beliefs through an – ever so slightly – violation of Bayes' rule. A naive equilibrium is a competitive equilibrium for a multiperiod complete markets economy such that every economic agent – Bayesian or non-Bayesian – assumes that all economic agents are Bayesian decision makers. If all agents are indeed Bayesian decision makers, the naive equilibrium coincides with the standard concept of an arbitrage-free equilibrium for which dynamic price ratios are comprehensively pinned down as the equilibrium price ratios of Arrow–Debreu securities in a static economy. If at least one agent is a non-Bayesian decision maker, however, some equilibrium price ratios will change over time. These changing price ratios imply the existence of unrealized dynamic arbitrage opportunities in a naive equilibrium with non-Bayesian decision makers.

1. Introduction

> "How do people form beliefs in situations under uncertainty? Economists have traditionally assumed that people begin with subjective beliefs over the different possible states of the world and then use Bayes' Rule to update those beliefs. This elegant and powerful model of economic agents as Bayesian statisticians is the foundation of modern information economics. Yet a large and growing body of psychological research suggests that the way people process information often departs systematically from Bayesian updating."Rabin and Schrag (1999), p.37)

Asset pricing models in mathematical finance are based on the paradigm that price processes must be arbitrage-free. For models with a finite state space this absence of arbitrage opportunities is equivalent to the existence of a state-price vector (i.e., a vector with prices for Arrow–Debreu securities) which allows to express asset prices as the assets' (discounted) expected payoffs with respect to a martingale measure. Asset pricing models in financial economics are based on competitive equilibria in asset exchange markets. Standard equilibrium models with expected utility maximizers characterize the equilibrium prices of Arrow–Debreu securities through the (i) the economic agents' risk-preferences as expressed by the curvature of their Bernoulli utility functions, (ii) the agents' beliefs given as additive probability measures that are updated in accordance with Bayes's rule, and (iii) the agents' initial endowments in Arrow–Debreu securities. Asset prices in these standard equilibrium models turn out to be arbitrage-free because the existence of any arbitrage opportunities would be incompatible with expected utility maximizing behavior of Bayesian decision makers.

This paper formally defines the concept of a 'naive equilibrium' for multiperiod complete markets economies with commonly observed information in which the economic agents can freely exchange Arrow–Debreu securities. At any possible information cell – i.e., at any date-event pair – the economic agents care about their expected utility from final consumption whereby this expectation is formed with respect to some information-conditional subjective probability measure. In contrast to the standard equilibrium concept, the economic agents in our model are not necessarily Bayesian decision makers. To be precise, recall that a Bayesian decision maker i is formally described through a *filtered probability space*

$$\left(\Omega, 2^{\Omega}, (\mathcal{F}_t)_{t\in\{0,\ldots,T\}}, \pi_i^0\right) \tag{1}$$

such that (i) Ω is a finite set which collect all economically relevant states, (ii) the sigma-algebra is given as the powerset 2^{Ω}, (iii) arrival of new information is governed by the commonly observed finite filtration process $(\mathcal{F}_t)_{t\in\{0,\ldots,T\}}$ consisting of σ-algebras generated by information partitions $(\Pi_t)_{t\in\{0,\ldots,T\}}$, and (iv) i's subjective prior belief is given as the additive probability measure π_i^0 on $\left(\Omega, 2^{\Omega}\right)$ which is updated in accordance with

☆ I would like to thank Alex Ludwig for helpful comments and suggestions. I am especially grateful to an anonymous reviewer and associate editor who both asked hard but good questions.

E-mail address: alexander.zimper@up.ac.za.

https://doi.org/10.1016/j.mathsocsci.2023.07.001
0165-4896/

Bayes' rule in the light of new information. To allow for non-Bayesian decision making, we extend (1) towards *a filtered belief process space*

$$\left(\Omega, 2^{\Omega}, (\mathcal{F}_t)_{t\in\{0,\ldots,T\}}, (\pi_{it})_{t\in\{0,\ldots,T\}}\right)$$

such that the function π_{it} assigns to every possible information cell $I_t \in \Pi_t$ in period t some subjective probability measure on I_t, denoted $\pi_i[I_t]$. As a generalization of (1), the posterior belief $\pi_i[I_t]$ is not necessarily derived from the prior π_i^0 through Bayes' rule.[1] We say that agent i is a non-Bayesian decision maker if and only if

$$\pi_i[I_t](\omega) \neq \frac{\pi_i^0(\{\omega\}\cap I_t)}{\pi_i^0(I_t)} = \pi_i^0(\omega \mid I_t)$$

for some information cell $I_t \in \Pi_t$ and some $\omega \in \Omega$.

In contrast to a Bayesian decision maker, the information conditional expectation of a non-Bayesian decision maker violates the *law of iterated expectations*, which is equivalent to *dynamic consistency*. The literature on dynamically inconsistent preferences distinguishes between the two extreme benchmark cases of *naive* versus *sophisticated* agents whereby empirical evidence suggests that naive decision making is a more prevalent phenomenon than sophisticated decision making (cf. O'Donoghue and Rabin (1999)). Whereas the sophisticated agent is fully aware of his dynamically inconsistent preferences, the naive agent incorrectly assumes that his preferences are dynamically consistent (see the literature cited in . Our concept of a naive equilibrium generalizes this notion of a naive decision maker from a single-agent to a multi-agent context: In a naive equilibrium every agent – generically incorrectly – assumes that all agents have dynamically consistent preferences, i.e., are Bayesian decision makers. As our main insight we establish equivalence between arbitrage-free equilibrium prices in a naive equilibrium and Bayesian decision making.

Main insight. *The following two assertions are equivalent for a naive equilibrium.*

(i) *The equilibrium price process is arbitrage-free.*

(ii) *All economic agents are Bayesian (i.e., dynamically consistent) decision makers.*

For analytical convenience, we show how the price process of any naive equilibrium can be equivalently generated through a 'representative agent' model. To be precise, we use the following notion of a representative agent.

> *We speak of a representative agent model for the price process of a naive equilibrium in the underlying multiperiod Arrow–Debreu economy iff (=if and only if) this price process is equivalently generated through a naive equilibrium in a single (=representative) agent multiperiod Arrow–Debreu economy such that this representative agent is – as the economic agents in the underlying economy – an expected utility maximizer.*

Because any equilibrium in a single-agent economy must be a zero-trade equilibrium, representative agent models come with the practical advantage that the naive equilibrium price ratios of the underlying economy are pinned down at all information cells through the first-order conditions of the representative agent's information-conditional expected utility maximization problem evaluated at his initial endowments. Reformulated for a representative agent economy our main insight becomes:

Main insight (representative agent version). *The following two assertions are equivalent for a naive equilibrium.*

(i) *The equilibrium price process is arbitrage-free.*

(ii) *The representative agent is a Bayesian (i.e., dynamically consistent) decision maker.*

The above results are sharp: marginal deviations from Bayesian updating by at least one economic agent must result in naive equilibrium prices that give rise to arbitrage opportunities. How plausible are deviations from Bayesian updating? Bayesian updating is a normatively important benchmark case because it ensures dynamically consistent expected utility preferences (cf., e.g., Epstein and Le Breton, 1993; Epstein and Schneider, 2003; Ghirardato, 2002; Siniscalchi, 2011). From a descriptive perspective, however, there exists a large body of psychological and economic literature which demonstrates that preferences of real life people are often subject to dynamic inconsistencies. In the specific context of processing new information, dynamic inconsistencies might arise whenever people are prone to psychological biases such as, e.g., overconfidence, confirmation, or/and myside biases.[2]

Whereas deviations from Bayesian updating seem to be the rule rather than the exception, it is less clear why existing arbitrage opportunities should not be exploited. The standard argument in favor of arbitrage-free asset prices is that some smart outside investor would quickly exploit any arbitrage opportunities which might arise in a naive equilibrium. In my opinion, the practical problem of exploiting arbitrage opportunities is not so much the lack of general knowledge that such opportunities exist but rather the lack of detailed knowledge about the right arbitrage trading strategy. For example, Hong and Stein (1999) argue that any investors who follow the news are simply too busy for understanding the specific arbitrage opportunities that might arise in a dynamic market context:

> "One can think of the newswatchers as having their hands full just figuring out the implications of the ϵ's for the terminal dividend D_T. This leaves them unable to [...] make any forecasts of future price changes, and hence unable to implement dynamic strategies [...]." (p.2149)

To illustrate the complexity of getting arbitrage strategies right, consider the following trading strategy: at first, trade Arrow–Debreu security s' for s in period zero; subsequently, reverse this trade in period t whenever information I_t is observed. For an economy in which all agents share the same CARA Bernoulli utility function, our representative agent analysis establishes that this trading strategy becomes an arbitrage strategy iff the n economic agents' non-Bayesian beliefs satisfy the following inequality

$$\prod_{i=1}^{n} \frac{\pi_i^0(\omega_s)}{\pi_i^0(\omega_{s'})} < \prod_{i=1}^{n} \frac{\pi_i[I_t](\omega_s)}{\pi_i[I_t](\omega_{s'})} \tag{2}$$

where $\pi_i^0(\omega)$ denotes agent i's prior probability and $\pi_i[I_t](\omega)$ denotes his information-conditional probability attached to state ω. Using this arbitrage strategy would be risky for an outside investor whenever he is not completely certain about (2): if this inequality was reversed, the investor would make sure losses

[1] Our general notion of a filtered belief process space includes the filtered probability space (1) of a Bayesian decision maker as a non-generic special case.

[2] The economic literature on non-Bayesian updating of additive probability measures is huge. Classic articles are Rabin and Schrag (1999), Rabin (2002), Epstein (2006), Epstein et al. (2008), Mullainathan et al. (2008), Gennaioli and Shleifer (2010), Ortoleva (2012). For more recent approaches see, e.g., the references in Baker (2022). Motivated by the empirical literature on 'asset mispricing', Daniel et al. (1998) model a confirmatory bias through non-Bayesian updating where a privately informed trader is overconfident with regards to his private information.

with the above trading strategy. To understand for which s and s' exactly inequality (2) holds, might be too difficult a task even for a quasi-smart investor who knows that inequality (2) must hold for some Arrow–Debreu securities s and s' because Bayesian updating is violated.

By construction, the naive equilibrium concept is only relevant for situations in which all agents – Bayesians and non-Bayesians – are naive in the specific sense that they are not aware of any dynamic inconsistencies in their own or the other agents' decision making. What would happen if some or all agents became sophisticated to the effect that they would correctly understand any dynamic inconsistencies in the model? In a single-agent decision theoretic context with dynamic inconsistencies a sophisticated agent would (i) either try to commit his future selves to his preferred plan of actions or (ii), in the absence of any viable commitment mechanism, the sophisticated agent would play a strategic game against all his future selves (see, e.g., Groneck et al. (2022) and references therein). In our multi-agent context, one would need to come up with a full-blown strategic model of strategic market interactions – with or without possibly costly commitment mechanisms – between a sophisticated agent and his futures selves as well with all other agents and their future selves. This is not a trivial task and beyond the scope of the present paper.

To come up with a proper game-theoretic model of sophisticated non-Bayesian decision makers who interact in a multiperiod Arrow–Debreu economy would be, in my opinion, a highly interesting – and highly relevant – task for future research. Clearly, sophisticated utility maximizing agents would exploit any arbitrage opportunities. While any subgame-perfect Nash equilibria of such market games with sophisticated agents must thus give rise to arbitrage-free asset prices, the resulting prices for non-Bayesian decision makers will be different than the arbitrage-free naive equilibrium prices for Bayesian decision makers.

The remainder of this paper is organized as follows. Section 2 discusses related literature. Section 3 introduces non-Bayesian belief processes. Section 4 defines the naive equilibrium concept for which we present representative agent models in Section 5. We show the equivalence between arbitrage-free naive equilibrium price processes and Bayesian updating in Section 6. The plausibility of unrealized arbitrage opportunities is discussed in Section 7. Section 8 concludes. Formal proofs are relegated to Appendix.

2. Relationship to the literature

2.1. Survival literature

The literature on dynamic survival models (cf. Sandroni (2000), Guerdjikova and Quiggin (2019), and references therein) argues that agents with incorrect beliefs may eventually run out of available resources because – in contrast to their correct belief counterparts – they tend to invest too much into objectively less likely states of the world. Blume and Easley (2006) combine this 'correct beliefs' argument with consistency results for Bayesian models of statistical learning (Doob, 1949; Berk, 1966). A Bayesian statistician holds a prior belief over a joint distributional parameter and data space whereby she uses Bayes' rule to form – in the light of observed data – posterior beliefs about the distributional parameter space. Because a Bayesian statistician will, in the limit, learn with probability one the true payoff distribution (provided that she considers the true distribution possible),[3] she will eventually outperform agents who are stuck with incorrect beliefs. During the review process I had been asked the following question by the associate editor:

> "Your model is one of complete markets and [...] you allow for the presence of some Bayesian agents. Could one view the arbitrage result as further rationalization of why only Bayesian traders asymptotically survive? If so, in what sense is the arbitrage result here more than a special case of the Blume and Easley survival result in the case where T is arbitrarily large (i.e., all non-Bayesian agents would eventually drop out because others profit at their expense)?"

The short answer is that our model has nothing to do with the question of whether a Bayesian decision maker outperforms a non-Bayesian decision maker or not. Our notion of a Bayesian decision maker is very different from the notion of a Bayesian statistician considered in Blume and Easley (2006). In contrast to a Bayesian statistician, the Bayesian decision maker of our model does not receive any statistical information but rather observes that some payoff-relevant states have become impossible. He then uses Bayes' rule to update his beliefs about the remaining payoff relevant states of the world, which simply means that his subjective posteriors coincide with the standard definition of conditional probability measures.[4] Since we do not consider any prior defined over some joint distributional parameter and sample space, there does not happen any Bayesian statistical learning in our model. Unlike in the dynamic survival literature, the beliefs of our Bayesian decision maker do thus not converge to any 'true' payoff distribution so that the 'converging to correct beliefs' argument by Blume and Easley (2006) does not apply.

To summarize: Unlike the dynamic survival literature, our model has nothing to say about whether Bayesian agents outperform non-Bayesian agents (or vice versa). We establish that unrealized (dynamic) arbitrage opportunities arise as an equilibrium phenomenon if (i) at least one agent is a non-Bayesian decision maker and (ii) all agents – Bayesian as well as non-Bayesian – are naive in the sense that they do not understand the arising arbitrage opportunities.

2.2. Mispricing and non-Bayesian updating literature

Our sharp characterization of arbitrage-free asset prices in terms of Bayesian updating comes with strong implications for the interpretation of empirically motivated 'asset mispricing' models.[5] 'Mispricing' in this literature often refers to the situation in which a risk-neutral representative agent prices an asset differently from its objective expected value because his belief differs from the objective payoff distribution. Mispricing in this specific sense thus already happens whenever the subjective beliefs of Bayesian decision makers do not coincide with the objective probability measure whereby such differences might even persist in the long run.[6] For an economy with (some) non-Bayesian

[3] If the true distribution does not belong to the support of the Bayesian statistician's prior, her Bayesian learning process will concentrate at distributions in her prior's support which minimize the Kullback and Leibler (1951) divergence (i.e., relative entropy) from the true distribution. For background reading see Zimper and Ma (2017) and references therein.

[4] For the according preference-based foundations of Bayesian updating in a conditional Savage (1954) world see, e.g., Epstein and Le Breton (1993), Ghirardato (2002), and – within a Choquet expected utility framework – Gilboa and Schmeidler (1993).

[5] For references to the large mispricing literature see, e.g., Ludwig and Zimper (2013).

[6] Barberis et al. (1998) formalize the empirical phenomena of over- versus underreaction of asset prices to news through a highly specific model of Bayesian updating. Their representative agent incorrectly believes that the payoff distribution is generated by either one of two Markov processes. Since the true payoff generating process – given as a random walk – does not belong to the support of the representative agent's prior, this Bayesian decison maker will never learn the true payoff distribution.

decision makers our analysis allows for a much stronger notion of 'mispricing', namely, the existence of unrealized arbitrage opportunities in a naive equilibrium.

The possibility of unrealized arbitrage opportunities appears to be neglected by existing asset pricing models which explicitly allow for non-Bayesian updating. As one example of this literature consider who write:

> "In our model, investors are quasi-rational in that they are Bayesian optimizers except for their overassessment of valid private information, and their biased updating of this precision." (p.1842)

Prices in these authors' model coincide with the subjective expectation of a privately informed (risk-neutral) trader who holds a subjective belief in the form of (i) a normally distributed prior and a (ii) data generating process that is independently normally distributed. Instead of using the 'correct formula' for the Bayesian posterior in this normal conjugate prior framework, the overconfident trader systematically overestimates the precision parameter of his private information. If this overconfident representative agent of Daniel et al. (1998) is indeed a non-Bayesian, i.e., dynamically inconsistent, decision maker in our sense – and not just a Bayesian decision maker whose subjective beliefs differ from objective probabilities –, then there would exist unrealized arbitrage opportunities in a complete markets version of their asset pricing model.

As another example, consider the heterogeneous beliefs model of Bhamra and Uppal (2014). These authors also explicitly allow for non-Bayesian updating:

> "[...] whereas it is possible to assume that beliefs are not updated at all, one could also assume Bayesian updating or some form of non-Bayesian updating." (Footnote 7, p.523)

Bhamra and Uppal (2014) use a de facto naive equilibrium framework:

> "We use a notion of equilibrium that is an extension of equilibrium in the single-agent model of Lucas (1978): both agents optimize their expected lifetime utility and all markets clear. Given our assumption that preferences are time separable and financial markets are complete, the dynamic consumption–portfolio choice problem simplifies to a static problem that requires one to choose the optimal allocation of consumption between the two investors for each date and state." (p.526)

Whereas our formal definition of a naive equilibrium restricts attention to final period consumption only, Bhamra and Uppal (2014) allow for intermediate consumption. The fact that we ignore—for the sake of analytical simplicity—any preferences for intermediate consumption—has no impact whatsoever on the characterization of arbitrage-free asset price-ratios whenever preferences are time separable as in Bhamra and Uppal (2014).[7] In contrast to Lucas's (1978) representative agent model with a Bayesian decision maker, the equilibrium prices in Bhamra and Uppal (2014) must therefore come with unrealized arbitrage opportunities whenever their economic agents are non-Bayesian decision makers.

Our argument concerning the pricing-models in Daniel et al. (1998) and in Bhamra and Uppal (2014) is as follows: while we consider these models as important and relevant, one needs to go one step further and emphasize the existence of unrealized arbitrage opportunities in a naive equilibrium whenever markets are complete and some non-Bayesian decision maker is present.

[7] We briefly sketch the according formal argument in Remark 3.

3. Belief processes

3.1. General set-up

Fix a finite state space $\Omega = \{\omega_1, \ldots, \omega_S\}$ with $S \geq 2$. Fix a finite number of time periods $t \in \{0, \ldots, T\}$ with $T \geq 1$. Denote by

$$\Pi_t = \left\{I_t^1, \ldots, I_t^{m_t}\right\}$$

a partition of the state space Ω into $m_t \geq 1$ information cells. In what follows we consider any fixed sequence of partitions $(\Pi_t)_{t\in\{0,\ldots,T\}}$ that satisfies the following two properties:

1. *$\Pi_0 = \{\Omega\}$ and $\Pi_T = \{\{\omega\} \mid \omega \in \Omega\}$;*
2. *Π_{t+1} is strictly finer than Π_t, that is, for every information cell $I_{t+1} \in \Pi_{t+1}$ there exists (i) some information cell $I_t \in \Pi_t$ such that $I_{t+1} \subseteq I_t$ whereby (ii) this inclusion is strict for at least one $I_t \in \Pi_t$.*

We write $I_t[\omega]$ to identify the unique period t information cell that contains state ω. Note that $I_t[\omega] \subseteq I_\tau[\omega]$ whenever $\tau < t$. $\mathcal{F}_t = \sigma(\Pi_t)$ denotes the sigma-algebra generated by Π_t whereby point 1. implies $\mathcal{F}_0 = \{\emptyset, \Omega\}$ and $\mathcal{F}_T = 2^\Omega$. $(\mathcal{F}_t)_{t\in\{0,\ldots,T\}}$ denotes the information filtration corresponding to $(\Pi_t)_{t\in\{0,\ldots,T\}}$. A sequence of mappings $\{y_t\}_{t\in\{0,\ldots,T\}}$ on Ω is an $(\mathcal{F}_t)_{t\in\{0,\ldots,T\}}$-adapted process iff for all $I_t \in \Pi_t$, $t \in \{0, \ldots, T\}$

$$y_t(\omega) = y_t\left(\omega'\right) \text{ whenever } \omega, \omega' \in I_t.$$

In that case we write $y[I_t(\omega)]$, or just $y[I_t]$, instead of $y_t(\omega)$ whenever it is understood that $\omega \in I_t$. We typically write y_0 for $y[I_0]$: since $y[I_0]$ takes on the same value for all $\omega \in I_0 = \Omega$, there is no ambiguity in doing so.

There are $n \geq 1$ different agents whereby each agent's uncertainty about the true state of the world is described by a subjective belief process that – somehow – incorporates the possible information that might be observed by the agent over time. Denote by π_{it} an $\mathcal{F}_t$ -measurable function on Ω that assigns to every $\omega \in \Omega$ some additive probability measure $\pi_i[I_t(\omega)]$ with full support on $\left(I_t(\omega), 2^{I_t}(\omega)\right)$. Again, instead of writing $\pi_{it}(\omega)$ for this probability measure we rather write $\pi_i[I_t(\omega)]$ or just $\pi_i[I_t]$ whenever it is understood that $\omega \in I_t$. For the corresponding probability of the state $\omega' \in I_t(\omega)$ we write $\pi_i[I_t(\omega)]\left(\omega'\right)$ or just $\pi_i[I_t]\left(\omega'\right)$. For the special case $t = 0$ the function π_{i0} is constant across all states because it assigns to every $\omega \in \Omega$ the same probability measure $\pi_i[I_0]$ on $\left(\Omega, 2^\Omega\right)$. We refer to $\pi_i[I_0]$ as the agent's prior belief. For notational simplicity we will typically write π_i^0 for the prior belief $\pi_i[I_0]$.

Definition 1. We call the $(\mathcal{F}_t)_{t\in\{0,\ldots,T\}}$-adapted process $(\pi_{it})_{t\in\{0,\ldots,T\}}$ with

$$\pi_{it} = \{\pi_i[I_t] \mid I_t \in \Pi_t\}$$

a belief process of agent i iff the following two conditions are satisfied for all $I_t \in \Pi_t$, $t \in \{0, \ldots, T\}$:

1. $\pi_i[I_t]$ is an additive probability measure defined on $\left(I_t, 2^{I_t}\right)$ with $\pi_i[I_t]\left(\omega'\right) > 0$ for all $\omega' \in I_t$;
2. $\pi_i[I_t] = \pi_i[I_{t-1}]$ whenever $I_t = I_{t-1}$.

3.2. Non-Bayesian versus Bayesian belief processes

Our notion of a Bayesian decision maker is adopted from the decision theoretic literature which investigates how subjective posterior beliefs are formed from subjective prior beliefs in a conditional Savage (1954) framework (cf. Epstein and Le Breton, 1993; Epstein and Schneider, 2003; Ghirardato, 2002). Under the assumption of a subjective expected utility maximizer, the updated posterior belief of a Bayesian decision maker is given as the (standard) conditional probability measure derived from his unconditional prior. Bayesian decision makers will be characterized through a Bayesian belief process whereas the belief processes of non-Bayesian decision makers violate Bayes' rule in some state of the world.

Definition 2. We call $(\pi_{it})_{t\in\{0,\dots,T\}}$ a Bayesian belief process if and only if, for all information cells $I_t \in \Pi_t$, $t \in \{1,\dots,T\}$, the posterior belief $\pi_i[I_t]$ is updated from the predecessor belief $\pi_i[I_{t-1}]$ with $I_{t-1} \supseteq I_t$ as a standard conditional probability measure in accordance with Bayes' rule, i.e.,

$$\pi_i[I_t](\omega) = \pi_i[I_{t-1}](\omega \mid I_t) = \frac{\pi_i[I_{t-1}](\omega)}{\pi_i[I_{t-1}](I_t)} \text{ for all } \omega \in I_t.$$

Repeating the argument

$$\frac{\pi_i[I_{t-1}](\omega)}{\pi_i[I_{t-1}](I_t)} = \frac{\frac{\pi_i[I_{t-2}](\omega)}{\pi_i[I_{t-2}](I_{t-1})}}{\frac{\pi_i[I_{t-2}](I_t)}{\pi_i[I_{t-2}](I_{t-1})}} = \frac{\pi_i[I_{t-2}](\omega)}{\pi_i[I_{t-2}](I_t)}$$

shows that the posterior beliefs $\pi_i[I_t]$, $I_t \in \Pi_t$, $t \in \{1,\dots,T\}$ of a Bayesian decision maker are equivalently given as conditional probability measures derived from his prior belief π_i^0 through Bayes' rule, i.e.,

$$\pi_i[I_t](\omega) = \pi_i^0(\omega \mid I_t) = \frac{\pi_i^0(\omega)}{\pi_i^0(I_t)} \text{ for all } \omega \in \Omega. \tag{3}$$

Fact 1. *While each agent i is associated with a filtered belief process space*

$$\left(\Omega, 2^{\Omega}, (\mathcal{F}_t)_{t\in\{0,\dots,T\}}, (\pi_{it})_{t\in\{0,\dots,T\}}\right), \tag{4}$$

this space (4) reduces to a filtered probability space

$$\left(\Omega, 2^{\Omega}, (\mathcal{F}_t)_{t\in\{0,\dots,T\}}, \pi_i^0\right)$$

satisfying (3) if and only if agent i is a Bayesian decision maker.

For a Bayesian decision maker we have, by (3), that

$$\begin{aligned} \mathbb{E}_{\pi_i[I_{t-1}]}Z &= \mathbb{E}_{\pi_i^0(\cdot\mid I_{t-1})}Z, \\ &\Leftrightarrow \\ \sum_{\omega\in I_{t-1}} Z(\omega)\,\pi_i[I_{t-1}](\omega) &= \sum_{\omega\in I_{t-1}} Z(\omega)\,\pi_i^0(\omega \mid I_{t-1}) \end{aligned}$$

as well as

$$\begin{aligned} \mathbb{E}_{\pi_i[I_{t-1}]}\left(\mathbb{E}_{\pi_i[I_t]}Z\right) &= \mathbb{E}_{\pi_i^0(\cdot\mid I_{t-1})}\left(\mathbb{E}_{\pi_i^0(\cdot\mid I_t)}Z\right) \\ &\Leftrightarrow \\ \sum_{I_t\subseteq I_{t-1}}\left(\sum_{\omega\in I_t} Z(\omega)\,\pi_i[I_t](\omega)\right)\pi_i[I_{t-1}](I_t) &= \sum_{I_t\subseteq I_{t-1}}\left(\sum_{\omega\in I_t} Z(\omega)\,\pi_i^0(\omega\mid I_t)\right) \\ &\quad \times \pi_i^0(I_t \mid I_{t-1}) \end{aligned}$$

for any $\mathcal{F}_T$-measurable random variable Z. By the law of iterated expectations

$$\mathbb{E}_{\pi_i^0(\cdot\mid I_{t-1})}\left(\mathbb{E}_{\pi_i^0(\cdot\mid I_t)}Z\right) = \mathbb{E}_{\pi_i^0(\cdot\mid I_{t-1})}Z \tag{5}$$

for the additive probability measure π_i^0, we thus obtain the following equivalent characterization of a Bayesian decision maker.

Fact 2. *Agent i a Bayesian decision maker if and only if his belief process $(\pi_{it})_{t\in\{0,\dots,T\}}$ satisfies the law of iterated expectations, that is, if and only if, for all $\mathcal{F}_T$-measurable random variables Z and all $I_t \in \Pi_t$, $t \in \{1,\dots,T\}$, and all $I_{t-1} \supseteq I_t$,*

$$\mathbb{E}_{\pi_i[I_{t-1}]}\left(\mathbb{E}_{\pi_i[I_t]}Z\right) = \mathbb{E}_{\pi_i[I_{t-1}]}Z. \tag{6}$$

Facts 1 and 2 establish that any economic agent i must be a Bayesian decision maker in our sense whenever his belief process is given as a filtered probability space

$$\left(\Omega, 2^{\Omega}, (\mathcal{F}_t)_{t\in\{0,\dots,T\}}, \pi_i^0\right) \tag{7}$$

such that the time-conditional expectation $\mathbb{E}_t^i Z$ of any $\mathcal{F}_T$-measurable random variable Z is an $\mathcal{F}_t$ -measurable random variable satisfies, for every $\omega \in I_t$,[8]

$$\begin{aligned} \mathbb{E}_t^i Z(\omega) = \mathbb{E}_{\pi_i^0(\cdot\mid I_t)}Z &= \int_{\omega'\in I_t} Z(\omega')\,d\pi_i^0(\cdot \mid I_t) \\ &= \frac{1}{\pi_i^0(I_t)}\int_{\omega'\in I_t} Z(\omega')\,d\pi_i^0. \end{aligned} \tag{8}$$

For our set-up of subjective expected utility maximizers, Bayesian decision making is also equivalent to dynamically consistent decision making (for a detailed formal argument in terms of preferences over Savage (1954) acts, see Epstein and Le Breton (1993))). This equivalence is also an immediate consequence of the equivalence between Bayesian decision making and the law of iterated expectations (6) as established by Fact 2. The basic (standard) argument goes as follows: under the law of iterated expectations, the maximization of information type I_{t-1}'s expected utility

$$\begin{aligned} \mathbb{E}_{\pi_i[I_{t-1}]}u_i(X,\theta_i) &= \mathbb{E}_{\pi_i[I_{t-1}]}\left(\mathbb{E}_{\pi_i[I_t]}u_i(X,\theta_i)\right) \\ &= \sum_{I_t\subseteq I_{t-1}} \pi_i[I_{t-1}](I_t) \\ &\quad \times \left(\sum_{\omega\in I_t} u_i(X(\omega),\theta_i[I_t])\,\pi_i[I_t](\omega)\right) \end{aligned} \tag{9}$$

over all $\theta_i[I_t]$, $I_t \subseteq I_{t-1}$, is equivalent to the maximization of each information type I_t's, $I_t \subseteq I_{t-1}$, expected utility

$$\mathbb{E}_{\pi_i[I_t]}u_i(X,\theta_i) = \left(\sum_{\omega\in I_t} u_i(X(\omega),\theta_i[I_t])\,\pi_i[I_t](\omega)\right)$$

over $\theta_i[I_t]$ because the probability $\pi_i[I_{t-1}](I_t)$ in (9) is a positive constant for each I_t. In other words, we have dynamic consistency because future information types $I_t \subseteq I_{t-1}$ have no incentive to deviate from an investment plan that is optimal from the perspective of the predecessor information type I_{t-1}.

Fact 3. *For subjective expected utility maximizers, Bayesian decision making is equivalent to dynamically consistent decision making.*

For the degenerate special case of a static economy, i.e., $T = 1$, all economic agents are trivially Bayesian decision makers. For $T \geq 2$ there exists, by construction, some information cell I_t with $\#I_t \geq 2$ such that the belief $\pi_i[I_t]$ can be any point in the interior of the $(\#I_t - 1)$-dimensional simplex whereas the Bayesian posterior $\pi_i[I_{t-1}](\cdot \mid I_t)$ is only a single point in this simplex. Since the Bayesian posterior has thus Lebesgue measure zero, a Bayesian decision maker corresponds to a non-generic special case of all possible belief processes whenever $T \geq 2$.

Example 1. Let $T = 3$ and consider a state space $\Omega = \{\omega_1,\dots,\omega_S\}$ with $S \geq 5$ and information partitions

$$\Pi_1 = \{\{\omega_1,\dots,\omega_4\},\{\omega_5,\dots,\omega_S\}\},$$

8 For a detailed analysis of the relationship between (i) conditional probability measures formed through Bayes' rule, (ii) conditional expectations operators, and (iii) the law of iterated expectations, see Section 34 in Billingsley (1995) (especially Example 34.1 and Theorem 34.4).

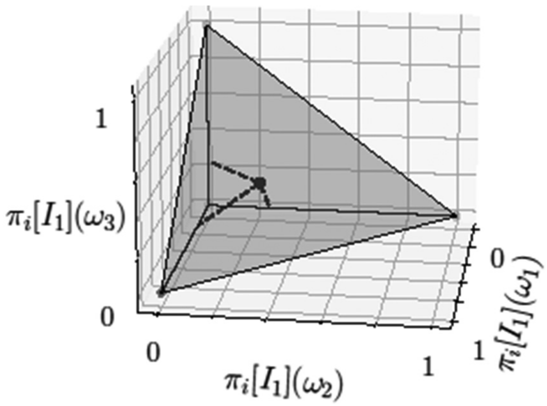

Fig. 1. The Bayesian posterior corresponds to the unique point $(\pi_i[I_1](\omega_1), \pi_i[I_1](\omega_2), \pi_i[I_1](\omega_3)) = \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right)$.

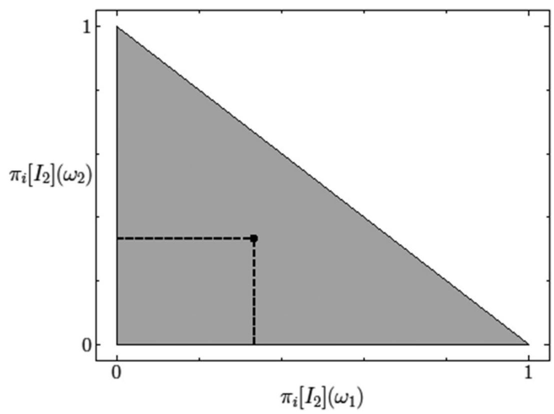

Fig. 2. The Bayesian posterior corresponds to the unique point $(\pi_i[I_2](\omega_1), \pi_i[I_2](\omega_2)) = \left(\frac{1}{3}, \frac{1}{3}\right)$.

$$\Pi_2 = \{\{\omega_1, \omega_2, \omega_3\}, \{\omega_4\}, \{\omega_5, \ldots, \omega_S\}\}.$$

Suppose that agent i has prior π_i^0 such that $\pi_i^0(\omega) = \frac{1}{S}$ for all $\omega \in \Omega$. If the agent is a Bayesian decision maker, he forms upon observing information $I_1 = \{\omega_1, \ldots, \omega_4\}$ and $I_2 = \{\omega_1, \omega_2, \omega_3\}$, respectively, the posteriors $\pi_i[I_1]$, resp. $\pi_i[I_2]$, such that

$$\pi_i[I_1](\omega) = \pi_i^0(\omega \mid I_1) = \begin{cases} \frac{1}{4} & \omega \in I_1 \\ 0 & else \end{cases}$$

$$\pi_i[I_2](\omega) = \pi_i^0(\omega \mid I_2) = \begin{cases} \frac{1}{3} & \omega \in I_2 \\ 0 & else. \end{cases}$$

These I_t-conditional probability measures correspond to the center of gravity $\left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right)$ for the 3 -dimensional (resp. $\left(\frac{1}{3}, \frac{1}{3}\right)$ for the 2 -dimensional) simplex (cf. Figs. 1 and 2). In contrast, a non-Bayesian agent i of our model can form any posterior belief $\pi_i[I_1]$ (resp. $\pi_i[I_2]$) that corresponds to an interior point in the 3-dimensional (resp. 2-dimensional) simplex.

Remark 1. It will not be relevant to our arbitrage analysis (i) whether there exists an objective (=true) probability measure on $\left(\Omega, 2^\Omega\right)$, denoted φ^0, or not, or (ii) whether the agents' subjective beliefs coincide with or converge to any existing objective measure φ^0. However, if we wanted to consider an agent i who has 'correct beliefs', we would model this agent through the filtered probability space

$$\left(\Omega, 2^\Omega, (\mathcal{F}_t)_{t\in\{0,\ldots,T\}}, \pi_i^0 = \varphi^0\right).$$

Note that an agent i with correct beliefs is thus a Bayesian decision maker whose posterior beliefs are given as conditional objective probability measures, i.e.,

$$\pi_i[I_t](\omega) = \pi_i^0(\omega \mid I_t) = \varphi^0(\omega \mid I_t) \text{ for all } \omega \in \Omega$$
$$\text{, all } I_t \in \Pi_t, t \geq 0.$$

4. Naive equilibria in a multiperiod Arrow–Debreu economy

We consider a complete markets multiperiod asset exchange economy in which S Arrow–Debreu securities can be traded between n agents in each pre-ultimate period $t \in \{0, \ldots, T-1\}$, $T \geq 2$. The Arrow–Debreu security $s \in \{1, \ldots, S\}$ pays out one unit of the consumption good in the ultimate period T if and only if the true state is ω_s. Formally, the Arrow–Debreu security s corresponds to the $\sigma(\Pi_T)$-measurable random variable $\mathbf{1}_{\{\omega_s\}}$, i.e., the indicator function of $\{\omega_s\}$. Note that Arrow–Debreu security s is worthless at any information cell I_t with $\omega_s \notin I_t$ because

$$\mathbf{1}_{\{\omega_s\}}(\omega) = 0 \text{ for all } \omega \in I_t.$$

For every agent i, fix some belief process $(\pi_{it})_{t\in\{0,\ldots,T\}}$ and some strictly increasing Bernoulli utility function $u_i : \mathbb{R}_{\geq 0} \to \{-\infty\} \cup \mathbb{R}$ defined over ultimate consumption in the final period T. Denote by $e_i^s[I_t]$ agent i's endowment of Arrow–Debreu security s at information cell I_t whenever $\omega_s \in I$. For simplicity, we assume that every agent initially owns some strictly positive endowment of every Arrow–Debreu security in period 0, i.e., $e_{i0} = e_i[I_0] \in \mathbb{R}^S_{>0}$ for $i \in \{1, \ldots, n\}$. At any given information cell $I_t \in \Pi_t, t \leq T-1$, agent i maximizes his expected utility over final period consumption with respect to the subjective belief $\pi_i[I_t]$.

Definition 3. Fix information cell $I_t \in \Pi_t$, $t \leq T-1$. Given the belief $\pi_i[I_t]$ and the I_t -endowments

$$e_i[I_t] = \left(e_i^s[I_t]\right)_{\{s|\omega_s\in I_t\}} \in \mathbb{R}^{\#I_t}_{\geq 0}$$

for all $i \in \{1, \ldots, n\}$, a static equilibrium at information cell I_t is an $(n+1)$ -tuple

$$\left(p^*[I_t], \theta_1^*[I_t], \ldots, \theta_n^*[I_t]\right) \in \mathbb{R}^{\#I_t}_{>0} \times \mathbb{R}^{\#I_t\times n}$$

that satisfies the following two conditions:

(i) Expected utility maximization at I_t : for $i \in \{1, \ldots, n\}$

$$\theta_i^*[I_t] \in \arg\max_{\theta_i\in B_i(p^*[I_t], e_i[I_t])} \sum_{\{s|\omega_s\in I_t\}} u_i\left(e_i^s[I_t] + \theta_i^s\right)\pi_i[I_t](\omega_s)$$

where

$$B_i\left(p^*[I_t], e_i[I_t]\right) = \left\{\theta_i \in \mathbb{R}^{\#I_t} \mid \sum_{\{s|\omega_s\in I_t\}} p^{*s}[I_t]\theta_i^s = 0 \text{ and } e_i^s[I_t] + \theta_i^s \geq 0 \text{ for all } s \text{ with } \omega_s \in I_t\right\}.$$

(ii) Market clearing at I_t :

$$\sum_{i=1}^n \theta_i^{*s}[I_t] = 0 \text{ for all } s \text{ with } \omega_s \in I_t.$$

We define a naive equilibrium as a collection of static equilibria for all information cells whereby we have to keep track of how initial endowments become, through successive net-trades, current endowments.

Definition 4. Fix the belief process $(\pi_{it})_{t\in\{0,\dots,T\}}$ and initial endowments $e_{i0} \in \mathbb{R}^S_{>0}$ for all agents $i \in \{1,\dots,n\}$. A naive equilibrium of the multiperiod Arrow–Debreu economy is an $(\mathcal{F}_t)_{t\in\{0,\dots,T\}}$-adapted process $\left(p_t^{**}, \theta_t^{**}\right)_{t\in\{0,\dots,T\}}$ defined as follows.

(i) Let $\omega_s \notin I_t$, i.e., Arrow–Debreu security s is worthless. We define

$$\left(p^{s**}[I_t], \theta_1^{s**}[I_t], \dots, \theta_n^{s**}[I_t]\right) = (0, 0, \dots, 0).$$

(ii) Let $\omega_s \in I_t$ such that $\{\omega_s\} \neq I_t$. We define

$$\left(p^{s**}[I_t], \theta_1^{s**}[I_t], \dots, \theta_n^{s**}[I_t]\right) = \left(p^{s*}[I_t], \theta_1^{s*}[I_t], \dots, \theta_n^{s*}[I_t]\right)$$

such that $\left(p^{s*}[I_t], \theta_1^{s*}[I_t], \dots, \theta_n^{s*}[I_t]\right)$ is a static equilibrium for the following I_t -endowments for all $i \in \{1,\dots, n\}$:

$$e_i^{s*}[I_t] = \begin{cases} e_{i0}^s & \text{if } t = 0 \\ e_{i0}^s + \sum_{\tau=0}^{t-1} \theta_i^{*s}[I_\tau] & \text{if } t > 0 \end{cases}. \tag{10}$$

(iii) Let $\omega_s \in I_t$ such that $\{\omega_s\} = I_t$. We define

$$p^{s**}[I_t] = 1,$$
$$\left(\theta_1^{s**}[I_t], \dots, \theta_n^{s**}[I_t]\right) = (0, \dots, 0).$$

A nave equilibrium is thus a collection of static equilibria for each information cell $I_t \in \Pi$ such that the $I_t(\omega)$ -endowments are handed down after trade at the predecessor information cell $I_{t-1}(\omega)$. If ω is the true state of the world, we would observe in a naive equilibrium $\left(p_t^{**}, \theta_t^{**}\right)_{t\in\{0,\dots,T\}}$ the following sequence of static equilibria over time:

$$\left(p^{s*}[I_t(\omega)], \theta_1^{s*}[I_t(\omega)], \dots, \theta_n^{s*}[I_t(\omega)]\right)_{t\in\{0,\dots,T\}}$$

for the corresponding endowments (10). Because a naive equilibrium exists if and only if the corresponding static equilibria exist, sufficiency conditions for the existence or/and uniqueness of naive equilibria are exactly the same as for static equilibria in complete market economies.[9] In what follows we focus on well-behaved naive equilibria for which any agent i's optimal net-trade decisions $\theta_i^{s**}[I_t] = \theta_i^{s*}[I_t]$ satisfy the first-order conditions

$$u_i'\left(e_i^{s*}[I_t] + \theta_i^{s*}[I_t]\right)\pi_i[I_t](\omega_s) = \lambda_i[I_t]\, p^{s**}[I_t] \text{ whenever } \omega_s \in I_t \tag{11}$$

where $\lambda_i[I_t]$ denotes agent i's Lagrange multiplier at information cell I_t.

Definition 5. We call $\left(p_t^{**}, \theta_t^{**}\right)_{t\in\{0,\dots,T\}}$ a 'well-behaved' naive equilibrium iff the equilibrium price-ratios satisfy, for all $I_t \in \Pi_t$, $t \in \{0,\dots,T-1\}$,

$$\frac{p^{s**}[I_t]}{p^{s'**}[I_t]} = \frac{u_i'\left(e_i^{s*}[I_t] + \theta_i^{s*}[I_t]\right)\pi_i[I_t](\omega_s)}{u_i'\left(e_i^{s'*}[I_t] + \theta_i^{s'*}[I_t]\right)\pi_i[I_t](\omega_{s'})} \text{ whenever } \omega_s, \omega_{s'} \in I_t \tag{12}$$

for an arbitrary agent $i \in \{1,\dots,n\}$.

Note that a naive equilibrium is always well-behaved if every agent's strictly increasing Bernoulli utility function is strictly concave and continuously differentiable on $(0,\infty)$. The next result is formally proved in Appendix.

[9] For existence and uniqueness results for competitive equilibria with additively separable utility functions, like the expected utility function of our model, see, e.g., Mas-Colell (1985), Dana (1993a,b)) and references therein.

Proposition 1. *Suppose that there exists a well-behaved naive equilibrium* $\left(p_t^{**}, \theta_t^{**}\right)_{t\in\{0,\dots,T\}}$*. Then the following two assertions are equivalent.*

(i) The price ratios are constant over time in the specific sense that for all $I_t \in \Pi_t$, $t \in \{0,\dots,T-1\}$

$$\frac{p^{s**}[I_t]}{p^{s'**}[I_t]} = \frac{p_0^{s*}}{p_0^{s'*}} = \frac{u_i'\left(e_{i0}^s + \theta_{i0}^{s*}\right)\pi_i^0(\omega_s)}{u_i'\left(e_{i0}^{s'} + \theta_{i0}^{s'*}\right)\pi_i^0(\omega_{s'})} \text{ whenever } \omega_s, \omega_{s'} \in I_t \tag{13}$$

for an arbitrary $i \in \{1,\dots,n\}$ *whereby* $\pi_i^0 = \pi_i[I_0]$ *denotes* i*'s prior belief on* $\left(\Omega, 2^\Omega\right)$.

(ii) All economic agents are Bayesian decision makers, i.e., for all $i \in \{1,\dots,n\}$ *and all* $I_t \in \Pi_t$, $t \in \{0,\dots,T-1\}$,

$$\pi_i[I_t](\omega) = \pi_i^0(\omega \mid I_t) = \frac{\pi_i^0(\omega)}{\pi_i^0(I_t)} \text{ for all } \omega \in I_t.$$

A naive equilibrium is a collection of standard equilibria such that all agents trade at any given information cell $I_t \in \Pi$ under the assumption that all agents are Bayesian decision makers. That is, every agent implicitly assumes at any given information cell $I_t \in \Pi$ that all agents $i \in \{1,\dots,n\}$ will use their current beliefs $\pi_i[I_t]$ as priors to derive all their future posteriors in accordance with Bayes' rule, i.e., for all $I_\tau \subseteq I_t$

$$\pi_i[I_\tau](\omega) = \pi_i[I_t](\omega \mid I_\tau) = \frac{\pi_i[I_t](\omega)}{\pi_i[I_t](I_\tau)} \text{ for all } \omega \in I_\tau.$$

If all agents are indeed Bayesian decision makers, any asset exchange in a naive equilibrium only happens in the first period. Afterwards, the agents have no further incentives to exchange Arrow–Debreu securities at the constant equilibrium price ratios (which are comprehensively pinned down as the price ratios of the static equilibrium at the initial information cell $I_0 = \Omega$). The naive equilibrium with Bayesian decision makers is therefore identical to the standard textbook notion of an equilibrium in a multiperiod complete-markets economy (cf., e.g., Chapter 4 in Arrow (1974), Chapter 19 in Mas-Colell et al. (1995), Chapter 2 in Duffie (2001)).

Remark 2. We restrict attention to Arrow–Debreu securities purely for analytical convenience. Our concept of a naive equilibrium naturally extends to multiperiod complete markets economies with arbitrary assets (i.e., random variables) under the standard portfolio-equivalence assumption that the equilibrium price $p^{X**}[I_t]$ of asset $X : \Omega \to \mathbb{R}$ at information cell $I_t \in \Pi$, $t \in \{0,\dots,T\}$, must coincide with the equilibrium value of the portfolio of Arrow–Debreu securities which replicates the payoffs of the random variable X on I_t. That is, in a naive equilibrium $\left(p_t^{**}, \theta_t^{**}\right)_{t\in\{0,\dots,T\}}$ the $(\mathcal{F}_t)_{t\in\{0,\dots,T\}}$-adapted naive equilibrium price process $\left(p_t^{X**}\right)_{t\in\{0,\dots,T\}}$ of asset X must satisfy at all $I_t \in \Pi$, $t \in \{0,\dots,T\}$

$$p^{X**}[I_t] = \sum_{\omega_s \in I_t} X(\omega_s)\, p^{s**}[I_t]. \tag{14}$$

For a static economy the asset pricing formula (14) in terms of the state price vector $\left(p_0^{s**}\right)_{s\in\{1,\dots,S\}}$ would guarantee that asset X is priced in an arbitrage-free way in period 0. The situation is different for multiperiod economies: whenever there is some agent who is not a Bayesian decision maker, the asset pricing formula (14) does no longer guarantee an arbitrage-free asset price for X because the naive equilibrium prices for Arrow–Debreu securities are themselves no longer arbitrage-free.

Remark 3. It does not matter for the analysis of equilibrium prices that our agents only care about final period consumption and not about intermediate consumption. To see this, consider the following slightly modified utility function at information cell I_t

$$u_i\left(c_i\left[I_t\right]\right)+\delta_{t,T}\sum_{\{s\mid\omega_s\in I_t\}}u_i\left(e_i^s\left[I_t\right]+\theta_i^s\right)\pi_i\left[I_t\right](\omega_s)$$

which (i) incorporates consumption at I_t given as

$$c_i\left[I_t\right]=\sum_{\{s\mid\omega_s\in I_t\}}p^s\left(e_i^s\left[I_t\right]-\theta_i^s\right)$$

and (ii) which discounts final period consumption by some time-discount factor $\delta_{t,T}>0$. In a naive equilibrium the first-order condition (11) becomes

$$u_i'\left(e_i^{s*}\left[I_t\right]+\theta_i^{s*}\left[I_t\right]\right)\pi_i\left[I_t\right](\omega_s)$$
$$=\delta_{t,T}^{-1}u_i'\left(c_i^*\left[I_t\right]\right)p^{s**}\left[I_t\right]\text{ whenever }\omega_s\in I_t.$$

Because $\delta_{t,T}^{-1}u_i'\left(c_i^*\left[I_t\right]\right)>0$ is a constant on I_t, the equilibrium prices ratios (13) remain unaffected for all $\omega_s,\omega_{s'}\in I_t$ by any intermediate consumption. Formally, agent i 's Lagrange multiplier $\lambda_i\left[I_t\right]$ in (11) is now his time adjusted marginal utility from (equilibrium) consumption at information cell I_t.

5. Representative agent models for naive equilibria

Since a naive equilibrium is a collection of static equilibria, the naive equilibrium price process can be conveniently described through a representative agent model for the static equilibria in question. Denote by $u_\rho:\mathbb{R}_{\geq0}\to\{-\infty\}\cup\mathbb{R}$ the strictly increasing Bernoulli utility function of the expected utility maximizing representative agent which is strictly concave and continuously differentiable on $(0,\infty)$. Furthermore, denote by $e_\rho=\left(e_\rho^1,\ldots,e_\rho^S\right)\in\mathbb{R}_{>0}^S$ the representative agent's initial endowments in Arrow–Debreu securities which remain constant over time.[10]

Definition 6. We say that there exists a representative agent model ρ for the naive equilibrium $\left(p_t^{**},\theta_t^{**}\right)_{t\in\{0,\ldots,T\}}$ iff we have for all $I_t\in\Pi_t$, $t\in\{0,\ldots,T-1\}$

$$\frac{p^{s**}\left[I_t\right]}{p^{s'**}\left[I_t\right]}=\frac{u_\rho'\left(e_\rho^s\right)}{u_\rho'\left(e_\rho^{s'}\right)}\frac{\pi_\rho\left[I_t\right](\omega_s)}{\pi_\rho\left[I_t\right](\omega_{s'})}\text{ whenever }\omega_s,\omega_{s'}\in I_t\tag{15}$$

where $\left(\pi_{\rho t}\right)_{t\in\{0,\ldots,T\}}$ is the $(\mathcal{F}_t)_{t\in\{0,\ldots,T\}}$-adapted belief process of the representative agent.

In words: the representative agent model ρ recovers the equilibrium prices of the underlying economy through an optimal zero net-trade decision of the expected utility maximizing representative agent at all information cells. The following result is analogously proved as Proposition 1.

Proposition 2. *If there exists a representative agent model ρ for the naive equilibrium $\left(p_t^{**},\theta_t^{**}\right)_{t\in\{0,\ldots,T\}}$, then the following two assertions are equivalent.*

(i) The equilibrium price ratios satisfy for all $I_t\in\Pi_t$, $t\in\{0,\ldots,T-1\}$

$$\frac{p^{s**}\left[I_t\right]}{p^{s'**}\left[I_t\right]}=\frac{u'\left(e_\rho^s\right)}{u'\left(e_\rho^{s'}\right)}\frac{\pi_{\rho0}(\omega_s)}{\pi_{\rho0}(\omega_{s'})}\text{ whenever }\omega_s,\omega_{s'}\in I_t$$

where $\pi_{\rho0}=\pi_\rho\left[I_0\right]$ denotes the representative agent's prior belief on $\left(\Omega,2^\Omega\right)$.

(ii) The representative agent is a Bayesian decision maker, that is, $\left(\pi_{\rho t}\right)_{t\in\{0,\ldots,T\}}$ satisfies, for all $I_t\in\Pi_t$, $t\in\{0,\ldots,T-1\}$,

$$\pi_\rho\left[I_t\right](\omega)=\pi_{\rho0}\left(\omega\mid I_t\right)=\frac{\pi_{\rho0}(\omega)}{\pi_{\rho0}\left(I_t\right)}\text{ for all }\omega\in I_t.$$

The next result – formally proved in the Appendix – establishes the existence of a representative agent model for any well-behaved naive equilibrium through the explicit construction of the corresponding aggregate belief process of the representative agent.

Theorem 1. *For every well-behaved naive equilibrium $\left(p_t^{**},\theta_t^{**}\right)_{t\in\{0,\ldots,T\}}$, there exists a representative agent model ρ such that the belief process $\left(\pi_{\rho t}\right)_{t\in\{0,\ldots,T\}}$ of the representative agent is characterized as follows: for all $I_t\in\Pi_t$, $t\in\{0,\ldots,T-1\}$,*

$$\pi_\rho\left[I_t\right](\omega_s)$$
$$=\frac{\frac{1}{u_\rho'\left(e_\rho^s\right)}\prod_{i=1}^n\left[u_i'\left(e_i^{s*}\left[I_t\right]+\theta_i^{s*}\left[I_t\right]\right)\pi_i\left[I_t\right](\omega_s)\right]^{\frac{\nu_i}{\sum_{i=1}^n\nu_i}}}{\sum_{\omega_{s'}\in I_t}\frac{1}{u_\rho'\left(e_\rho^{s'}\right)}\prod_{i=1}^n\left[u_i'\left(e_i^{s'*}\left[I_t\right]+\theta_i^{s'*}\left[I_t\right]\right)\pi_i\left[I_t\right](\omega_{s'})\right]^{\frac{\nu_i}{\sum_{i=1}^n\nu_i}}}$$
$$\text{for all }\omega_s\in I_t\tag{16}$$

where the agent coefficients $\nu_i>0$, $i\in\{1,\ldots,n\}$, are arbitrary.

The remainder of this section uses Theorem 1 to derive analytically convenient representative agent models for CARA and CRRA Bernoulli utility functions, respectively.

Remark 4. In Zimper (2023) I study representative agent models for equilibria in static complete markets economies. There I derive belief aggregation formulas under the assumption that all economic agents share the same well-behaved Bernoulli utility functions. For the special case of CARA and CRRA Bernoulli utility functions the corresponding representative agent models are mathematically equivalent to existing aggregation results in Rubinstein (1974, 1976), Jouini and Napp (2007), Calvet et al. (2018).[11] The main difference of my approach and this existing literature is basically a normalization exercise whose advantage is of expositional nature: whereas the present paper and Zimper (2023) require representative agents to be expected utility maximizers with aggregate beliefs given as additive probability measures, Jouini and Napp (2007) use a – mathematically equivalent – 'consensus characteristic' approach which is only an additive probability measure for the special case of logarithmic CRRA Bernoulli utility. To describe the aggregate beliefs of representative agents as additive probability measures comes with the important conceptual advantage that one can meaningfully distinguish between Bayesian versus non-Bayesian expected utility maximizing representative agents.

5.1. CARA Bernoulli utility

Assume that every agent has some Bernoulli utility function which is of the CARA form, that is, $u_i:\mathbb{R}_{\geq0}\to\{-\infty\}\cup\mathbb{R}_{<0}$ such

10 While it is standard in the literature to consider either average, i.e., $e_\rho^s=\frac{1}{n}\sum_{i=1}^n e_i^s$, or aggregate, i.e., $e_\rho^s=\sum_{i=1}^n e_i^s$, endowments for all $s\in\{1,\ldots,S\}$, the present paper allows for more general endowment specifications for the representative agent (cf. Propositions 3 and 4).

11 For a continuum of economic agents on the real line with CRRA Bernoulli utility function, Atmaz and Basak (2018) derive an analogous representative agent result for a diffusion process under the assumption that the different beliefs across the agents about the mean of the dividend growth rate are normally distributed around the true mean. Similarly, for a continuum of economic agents on the open unit interval with logarithmic CRRA Bernoulli utility function, Martin and Papadimitriou (2021) present a representative agent result for a binomial tree process.

that

$$u_i(c) = \begin{cases} -\exp(-\alpha_i c) & \text{if } c > 0 \\ \infty & \text{if } c = 0 \end{cases}$$

where $\alpha_i > 0$ stands for agent i's absolute risk aversion coefficient. We prove the following result in Appendix.

Proposition 3. *Suppose that all economic agents have CARA Bernoulli utility functions with $\alpha_i > 0$ for all i. Then there exists a representative agent model ρ for the well-behaved naive equilibrium $\left(p_t^{**}, \theta_t^{**}\right)_{t\in\{0,\dots,T\}}$ described as follows:*

(i) The Bernoulli utility function u_ρ of the representative agent is also of the CARA form whereby we can fix an arbitrary CARA coefficient $\alpha_\rho > 0$.
(ii) The initial endowments of the representative agent satisfy

$$e_\rho^s = \frac{1}{\sum_{i=1}^n \frac{\alpha_\rho}{\alpha_i}} \sum_{i=1}^n e_{i0}^s \text{ for all } s \in \{1,\dots,S\}.$$

(iii) The belief process $\left(\pi_{\rho t}\right)_{t\in\{0,\dots,T\}}$ of the representative agent is characterized as follows: for all $I_t \in \Pi_t$, $t \in \{0,\dots,T-1\}$,

$$\pi_\rho[I_t](\omega_s) = \frac{\left(\prod_{i=1}^n [\pi_i[I_t](\omega_s)]^{\frac{\alpha_\rho}{\alpha_i}}\right)^{\frac{1}{\sum_{i=1}\frac{\alpha_\rho}{\alpha_i}}}}{\sum_{\omega_{s'}\in I_t}\left(\prod_{i=1}^n [\pi_i[I_t](\omega_{s'})]^{\frac{\alpha_\rho}{\alpha_i}}\right)^{\frac{1}{\sum_{i=1}\frac{\alpha_\rho}{\alpha_i}}}} \text{ for all } \omega_s \in I_t.$$

Corollary 1. *Suppose that all agents share the same CARA Bernoulli utility function, i.e., $\alpha = \alpha_\rho = \alpha_i$ for all i. Then the representative agent model of Proposition 3 becomes an average endowment model, i.e.,*

$$e_\rho^s = \frac{1}{n}\sum_{i=1}^n e_{i0}^s \text{ for all } s \in \{1,\dots,S\},$$

such that the aggregate beliefs are given as

$$\pi_\rho[I_t](\omega_s) = \frac{\left(\prod_{i=1}^n \pi_i[I_t](\omega_s)\right)^{\frac{1}{n}}}{\sum_{\omega_{s'}\in I_t}\left(\prod_{i=1}^n \pi_i[I_t](\omega_{s'})\right)^{\frac{1}{n}}} \text{ for all } \omega_s \in I_t. \tag{17}$$

Moreover, if all economic agents additionally share the same belief, i.e., $\pi[I_t] = \pi_i[I_t]$ for all i, then (17) simplifies to $\pi_\rho[I_t] = \pi[I_t]$.

5.2. CRRA Bernoulli utility

Next assume that all agents share the same CRRA Bernoulli utility function $u : \mathbb{R}_{\geq 0} \to \{-\infty\} \cup \mathbb{R}$ such that, for $c > 0$,

$$u(c) = \begin{cases} \frac{c^{1-\gamma}}{1-\gamma} & \text{if } \gamma \neq 1 \\ \ln c & \text{if } \gamma = 1 \end{cases}$$

and

$$u(0) = \begin{cases} \frac{c^{1-\gamma}}{1-\gamma} & \text{if } \gamma < 1 \\ -\infty & \text{if } \gamma \geq 1 \end{cases}$$

where $\gamma > 0$ stands for the relative risk aversion coefficient shared by all agents. Recall from (11) that agent i's Lagrange multiplier at information cell I_t evaluated at equilibrium values is given as

$$\lambda_i[I_t] = \frac{1}{p^{s**}[I_t]} u'\left(e_i^{s*}[I_t] + \theta_i^{s*}[I_t]\right)\pi_i[I_t](\omega_s). \tag{18}$$

The next result is formally proved in Appendix.

Proposition 4. *Suppose that all economic agents share the same CRRA Bernoulli utility function. Then there exists a representative agent model ρ for the well-behaved naive equilibrium $\left(p_t^{**}, \theta_t^{**}\right)_{t\in\{0,\dots,T\}}$ described as follows:*

(i) The Bernoulli utility function u_ρ of the representative agent is of the same CRRA form with $\gamma > 0$.
(ii) The initial endowments of the representative agent satisfy

$$e_\rho^s = a\sum_{i=1}^n e_{i0}^s \text{ for all } s \in \{1,\dots,S\}$$

for an arbitrary $a > 0$.
(iii) The belief process $\left(\pi_{\rho t}\right)_{t\in\{0,\dots,T\}}$ of the representative agent is characterized as follows: for all $I_t \in \Pi_t$, $t \in \{0,\dots,T-1\}$,

$$\pi_\rho[I_t](\omega_s) = \frac{\left(\sum_{i=1}^n \left(\hat{\mu}_i[I_t]\pi_i[I_t](\omega_s)\right)^{\frac{1}{\gamma}}\right)^\gamma}{\sum_{\omega_{s'}\in I_t}\left(\sum_{i=1}^n \left(\hat{\mu}_i[I_t]\pi_i[I_t](\omega_{s'})\right)^{\frac{1}{\gamma}}\right)^\gamma} \text{ for all } \omega_s \in I_t \tag{19}$$

whereby the agent weights are given as

$$\hat{\mu}_i[I_t] = \frac{(\lambda_i[I_t])^{-1}}{b} \tag{20}$$

for an arbitrary constant $b > 0$ with $\lambda_i[I_t]$ denoting agent i's Lagrange multiplier (18).

In contrast to the CARA representative agent model, the CRRA representative agent model with heterogeneous beliefs has to keep track of the economic agents' equilibrium trades because the past and present equilibrium trades enter into each agent's Lagrange multiplier (18), which in turn enters into the agent weight (20). Whenever the agents have heterogeneous beliefs it is therefore analytically much more convenient to work with CARA instead of CRRA Bernoulli utility. This problem does not arise for the CRRA representative agent model if the economic agents have homogeneous beliefs: note that

$$\frac{\left(\sum_{i=1}^n \left(\hat{\mu}_i[I_t]\pi[I_t](\omega_s)\right)^{\frac{1}{\gamma}}\right)^\gamma}{\sum_{\omega_{s'}\in I_t}\left(\sum_{i=1}^n \left(\hat{\mu}_i[I_t]\pi[I_t](\omega_{s'})\right)^{\frac{1}{\gamma}}\right)^\gamma} = \frac{\pi[I_t](\omega)\left(\sum_{i=1}^n \left(\hat{\mu}_i[I_t]\right)^{\frac{1}{\gamma}}\right)^\gamma}{\sum_{\omega_{s'}\in I_t}\pi[I_t](\omega_{s'})\left(\sum_{i=1}^n \left(\hat{\mu}_i[I_t]\right)^{\frac{1}{\gamma}}\right)^\gamma} = \pi[I_t](\omega)$$

because of $\sum_{\omega_{s'}\in I_t}\pi[I_t](\omega_{s'}) = 1$. This gives us the following corollary.

Corollary 2. *Suppose that all economic agents share the same belief, i.e., $\pi[I_t] = \pi_i[I_t]$ for all i. Then (19) simplifies to $\pi_\rho[I_t] = \pi[I_t]$.*

Remark 5. Since we are free to choose $b > 0$ in (20), we can set

$$b = \sum_{j=1}^n \left(\lambda_j[I_t]\right)^{-1}$$

to obtain the normalized agent weights

$$\mu_i^*[I_t] = \frac{(\lambda_i[I_t])^{-1}}{\sum_{j=1}^n \left(\lambda_j[I_t]\right)^{-1}}, \tag{21}$$

which sum up to one. The analysis of the static economy in the companion paper Zimper (2023) shows that these normalized agent weights are 'Pareto' agent weights in the specific sense that the utilitarian welfare function

$$W_{\mu^*}\left(\theta_1^s[I_t],\ldots,\theta_n^s[I_t]\right) = \sum_{i=1}^n \mu_i^* E_{\pi_i[I_t]} u\left(\theta_i^s[I_t]+e_i^{s*}[I_t]\right)$$
$$= \sum_{i=1}^n \mu_i^* \sum_{\omega_s\in I_t} u\left(\theta_i^s[I_t]+e_i^{s*}[I_t]\right)\pi_i(\omega_s)$$
$$\text{subject to} \sum_{i=1}^n \theta_i^s[I_t] = 0 \text{ and } \theta_i^s[I_t]+e_i^s[I_t]\geq 0 \text{ for all } \omega_s\in I_t$$

is maximized at the naive equilibrium net-trade decisions $\left(\theta_1^{s*}[I_t],\ldots,\theta_n^{s*}[I_t]\right)$. This equivalence between equilibrium net-trades and utilitarian welfare maximization with appropriately constructed agent weights directly confirms the well-known fact that competitive equilibria in complete markets economies are Pareto-efficient.[12]

6. Naive equilibria versus arbitrage-free asset prices

This section establishes that arbitrage-free prices in a naive equilibrium are equivalent to Bayesian decision making. Denote by $(p_t)_{t\in\{0,\ldots,T\}}$ an $(\mathcal{F}_t)_{t\in\{0,\ldots,T\}}$-adapted price process satisfying

1. $p^s[I_t]>0$ if and only if $\omega_s\in I_t$.
2. $p^s[I_T]=\mathbf{1}_{\{\omega_s\}}$.

For a fixed price process $(p_t)_{t\in\{0,\ldots,T\}}$ the $(\mathcal{F}_t)_{t\in\{0,\ldots,T\}}$-adapted trading strategy $\{\theta_t\}_{t\in\{0,\ldots,T\}}$ is an *arbitrage-opportunity* in our multiperiod Arrow–Debreu economy iff the following two conditions are satisfied:

1. It is self-financing, i.e., for all $I_t\in\Pi_t$, $t\in\{0,\ldots,T-1\}$,
$$\sum_{\{s\mid\omega_s\in I_t\}} p^s[I_t]\,\theta^s[I_t]=0;$$
2. The resulting period $T-1$ portfolio strictly dominates the initial endowment portfolio in the specific sense that, for all $s\in\{1,\ldots,S\}$,
$$\sum_{t=0}^{T-1}\theta^s[I_t(\omega_s)]\geq 0$$
whereby this inequality is strict for some s.

We speak of an arbitrage-free price process iff there is no arbitrage-opportunity. By a fundamental result from mathematical finance, the price process $(p_t)_{t\in\{0,\ldots,T\}}$ is arbitrage-free iff it is a (discounted) martingale with respect to an additive probability measure Q with full support on Ω; that is, iff we have for all $s\in\{1,\ldots,S\}$ and all $I_t\in\Pi_t$, $t\in\{1,\ldots,T-1\}$,

$$p^s[I_t]=\frac{1}{1+r_{t,t+1}[I_t]}\mathbb{E}_{Q(\cdot|I_t)}p^s[I_{t+1}] \quad (22)$$

where $Q(\omega)>0$ for all ω and $\left(r_{t,t+1}\right)_{t\in\{0,\ldots,T\}}$ is an arbitrary $(\mathcal{F}_t)_{t\in\{0,\ldots,T\}}$-adapted short-rate process satisfying, for all $\omega\in\Omega$, $r_{T,T+1}[I_T(\omega)]=0$ and $1+r_{t,t+1}[I_t(\omega)]>0$ for $t<T$ (cf., Harrison and Kreps (1979), Chapter 2G in Duffie (2001)). Moreover, for our complete markets economy the martingale measure Q is unique.

[12] Traditional aggregation analysis has used this Pareto-efficiency of equilibria to derive representative agent models from the maximization of the utilitarian welfare function W_{μ^*} whereby the agent weights $\mu_i^*[I_t]$, $i\in\{1,\ldots,n\}$, are given by equivalent expressions of (21). See, e.g., Negishi (1960), the discussion in Zimper (2023), and, in particular, Chapter 1E on representative agent models in Duffie (2001) who writes: "Aside from its allocational implications, Pareto optimality is also a convenient property for the purpose of security pricing."(p.8)

We formally prove the following result in the Appendix through the explicit construction of the martingale measure Q.

Proposition 5. *Fix a period-0 state-price vector*

$$p_0=\left(p_0^1,\ldots,p_0^S\right)\in\mathbb{R}^S_{>0}. \quad (23)$$

The $(\mathcal{F}_t)_{t\in\{0,\ldots,T\}}$-adapted price process $(p_t)_{t\in\{0,\ldots,T\}}$ is arbitrage free iff we have for all $I_t\in\Pi_t$, $t\in\{1,\ldots,T-1\}$,

$$\frac{p_0^s}{p_0^{s'}}=\frac{p^s[I_t]}{p^{s'}[I_t]} \text{ whenever } \omega_s,\omega_{s'}\in I_t. \quad (24)$$

Combining Propositions 1 and 2 with Proposition 5 yields the following fundamental relationship.

Theorem 2. *Consider a well-behaved naive equilibrium $\left(p_t^{**},\theta_t^{**}\right)_{t\in\{0,\ldots,T\}}$. The following assertions are equivalent.*

(i) The equilibrium price process $\left(p_t^{**}\right)_{t\in\{0,\ldots,T\}}$ is arbitrage-free.
(ii) All economic agents are Bayesian decision makers.
(iii) The representative agent is a Bayesian decision maker.

Remark 6. The $(\mathcal{F}_t)_{t\in\{0,\ldots,T\}}$-adapted short-rate process $\left(r_{t,t+1}\right)_{t\in\{0,\ldots,T\}}$ is an arbitrary normalization/numeraire process which is not pinned down in our economy because our agents only care about final period consumption. Define the $\sigma(\Pi_t)$-measurable random variable $R_{t,T}$ such that for all $I_t\in\Pi_t$

$$R_{t,T}[I_t]=\prod_{\tau=t}^T\left(1+r_{\tau,\tau+1}[I_\tau]\right).$$

By (22) and the law of iterated expectations, we obtain

$$p^s[I_t]=\frac{1}{R_{t,T}[I_t]}\mathbb{E}_{Q(\cdot|I_t)}p^s[I_T]=\frac{1}{R_{t,T}[I_t]}Q(\omega_s\mid I_t).$$

At any information cell I_t the risk-free asset (i.e., portfolio of Arrow–Debreu securities) that gives a guaranteed payoff of one in period T would thus come with the price

$$\sum_{\omega_s\in I_t}p^s[I_t]=\sum_{\omega_s\in I_t}\frac{1}{R_{t,T}[I_t]}Q(\omega_s\mid I_t)$$
$$=\frac{1}{R_{t,T}[I_t]}.$$

In other words, the short-rate process $\left\{r_{t,t+1}\right\}_{t\in\{0,\ldots,T\}}$ determines the nominal price $\frac{1}{R_{t,T}[I_t]}$ that has to be paid at information cell I_t for one unit of a risk-free asset whose period-T payoff is one.

7. Discussion: Unrealized arbitrage opportunities in a naive equilibrium

How plausible is the existence of unrealized arbitrage opportunities? By Proposition 5, we obtain the following equivalent characterization of arbitrage opportunities in a naive equilibrium.

Corollary 3. *The equilibrium price process $\left(p_t^{**}\right)_{t\in\{0,\ldots,T\}}$ comes with unrealized arbitrage opportunities iff there exist some Arrow–Debreu securities s,s' such that*

$$\frac{p_0^{s**}}{p_0^{s'**}}<\frac{p^{s**}[I_t]}{p^{s'**}[I_t]} \quad (25)$$

for some $I_t\in\Pi_t$ with $\omega_s,\omega_{s'}\in I_t$.

Suppose that there exists a smart investor who knows in period 0 that inequality (25) holds for fixed Arrow–Debreu securities s and s'. This investor could then use an arbitrage strategy according to which he sells in period 0 (as many as possible) units of Arrow–Debreu security s' in order to buy from the proceeds units of Arrow–Debreu security s whereas this trade is reversed whenever information cell I_t happens. We refer to this arbitrage strategy as AS^*. More precisely, arbitrage strategy AS^* is given as the self-financing trading strategy[13]

$$\begin{aligned}\theta_0^{s'} &< 0,\\ \theta_0^{s} &= -\frac{p_0^{s'**}}{p_0^{s**}}\theta_0^{s'},\\ \theta^{s}[I_t] &= -\theta_0^{s},\\ \theta^{s'}[I_t] &= -\frac{p^{s**}[I_t]}{p^{s'**}[I_t]}\theta^{s}[I_t]\end{aligned}$$

which results in the following strictly dominating portfolio

$$\begin{aligned}\theta_0^{s}+\theta^{s}[I_t] &= 0,\\ \theta_0^{s'}+\theta^{s'}[I_t] &= \left(1-\frac{p^{s**}[I_t]}{p^{s'**}[I_t]}\frac{p_0^{s'**}}{p_0^{s**}}\right)\theta_0^{s'} > 0\end{aligned}$$

whereby the last line follows because (25) implies

$$1-\frac{p^{s**}[I_t]}{p^{s'**}[I_t]}\frac{p_0^{s'**}}{p_0^{s**}} < 0.$$

That is, if information cell I_t happens, the initial trade combined with the reversed trade in Arrow–Debreu securities s and s' at information cell I_t results in a portfolio with strictly more units of the Arrow–Debreu security s' than in period 0 whereas the number of units of Arrow–Debreu security s remains the same. If information cell I_t does not happen in period t because the true state of the world lies in a different period t information cell, both Arrow–Debreu securities s and s' are worthless anyway so that the initial trade does not affect the value of the portfolio in period t.

How can this smart investor know in period 0 that inequality (25) holds for fixed Arrow–Debreu securities s and s'? A 'fully rational' investor knows, by definition, all economic fundamentals, i.e., all agents' endowments, utility functions and belief processes, from which he could deduce inequality (25) if he additionally understands that the economic agents trade in accordance with a naive equilibrium. Let us simplify the task of this investor by assuming that he happens to know that all economic agents share the same CARA Bernoulli utility function. We can then use the belief aggregation formula for the CARA representative agent model of Corollary 1 to obtain the following result.

Corollary 4. *Suppose that all economic agents $i \in \{1, \ldots, n\}$ share the same CARA Bernoulli utility function. Then inequality (25) holds iff*

$$\prod_{i=1}^{n}\frac{\pi_i^0(\omega_s)}{\pi_i^0(\omega_{s'})} < \prod_{i=1}^{n}\frac{\pi_i[I_t](\omega_s)}{\pi_i[I_t](\omega_{s'})}. \tag{26}$$

Consequently, whenever the smart investor additionally knows that the economic agents' beliefs satisfy inequality (26) for fixed Arrow–Debreu securities s and s', he also knows that the arbitrage strategy AS^* would strictly improve his portfolio at zero costs. If this smart investor had sufficient endowments to make a difference, the naive equilibrium concept would not be an adequate model for asset prices. However, it is hard to imagine how real-life investors could possibly know that inequality (26) holds for fixed Arrow–Debreu securities s and s'.

To conclude: Even if a quasi-smart outside investor knows that there must be some arbitrage opportunities because not all economic agents are Bayesian decision makers, it might be difficult for him to figure out the correct arbitrage trading strategy. If this investor accidentally happens to mix up the two Arrow–Debreu securities s and s' in (26), he would make strict losses at information cell I_t (and zero-gains elsewhere) whenever he uses the trading strategy AS^*.

8. Concluding remarks

This paper introduces the concept of 'naive equilibria' for complete markets multiperiod economies with Arrow–Debreu securities whereby the economic agents are not necessarily Bayesian decision makers. A naive equilibrium is the adequate equilibrium concept if every economic agent – generically incorrectly – assumes that all agents are Bayesian decision makers. For the non-generic case where all agents are indeed Bayesian decision makers, the naive equilibrium coincides with the standard concept of an arbitrage-free equilibrium in a multiperiod complete markets economy. In such a standard arbitrage-free equilibrium dynamic price ratios are comprehensively pinned down as the equilibrium price ratios of a static economy and the economic agents have no strict incentives to trade Arrow–Debreu securities beyond the initial trading period. Our analysis shows that the situation is different for a naive equilibrium if there exists at least one agent who is a non-Bayesian decision maker. In this generic case some equilibrium price ratios for Arrow–Debreu securities will change over time. These changing price ratios imply the existence of unrealized dynamic arbitrage opportunities in a naive equilibrium.

Data availability

No data was used for the research described in the article.

Appendix. Mathematical proofs

Proof of Proposition 1. Part (i). Suppose that all agents are Bayesian decision makers. Then the two maximization problems

$$c_i^*[I_t] \in \arg\max_{c_i\in B_i(p^*[I_t])}\frac{1}{\pi_i^0(I_t)}\sum_{\{s|\omega_s\in I_t\}}u_i\left(c_i^s\right)\pi_i^0(\omega_s)$$

and

$$c_{i0}^* \in \arg\max_{c_i\in B_i(p_0^*)}\sum_{\{s|\omega_s\in\Omega\}}u_i\left(c_i^s\right)\pi_i^0(\omega_s)$$

identically yield for all $\omega_s \in I_t$

$$\begin{aligned}c_i^{*s}[I_t] &= c_{i0}^{*s}\\ &\Leftrightarrow\\ e_i^{s*}[I_t]+\theta_i^{s*}[I_t] &= e_{i0}^{s}+\theta_{i0}^{s*}\end{aligned}$$

for all $I_t \in \Pi_t$, $t \in \{0, \ldots, T-1\}$. By (12), we have for all $I_t \in \Pi_t$, $t \in \{0, \ldots, T-1\}$,

$$\frac{p^{s**}[I_t]}{p^{s'**}[I_t]} = \frac{u_i'\left(e_{i0}^{s}+\theta_{i0}^{s*}\right)\pi_i^0(\omega_s)}{u_i'\left(e_{i0}^{s'}+\theta_{i0}^{s'*}\right)\pi_i^0(\omega_{s'})} \text{ whenever } \omega_s, \omega_{s'} \in I_t.$$

Since the right hand side of the equation has the same value for all I_t with $\omega_s, \omega_{s'} \in I_t$, we obtain (13).

[13] Let $e_0^{s'} > 0$ denote the initial endowment of the smart investor of Arrow–Debreu security s'. We can then set $\theta_0^{s'} = -e_0^{s'}$ so that the smart investor would maximally exploit this arbitrage opportunity in the absence of naked short-selling.

$$\frac{\prod_{i=1}^{n}\left[u_i'\left(e_i^{s*}[I_t]+\theta_i^{s*}[I_t]\right)\pi_i[I_t](\omega_s)\right]^{\nu_i}}{\prod_{i=1}^{n}\left[u_i'\left(e_i^{s'*}[I_t]+\theta_i^{s'*}[I_t]\right)\pi_i[I_t](\omega_{s'})\right]^{\nu_i}}=\left(\frac{u_\rho'\left(e_\rho^{s}\right)}{u_\rho'\left(e_\rho^{s'}\right)}\frac{\pi_\rho[I_t](\omega_s)}{\pi_\rho[I_t](\omega_{s'})}\right)^{\sum_{i=1}^{n}\nu_i}$$

$$\Leftrightarrow$$

$$\frac{\pi_\rho[I_t](\omega_s)}{\pi_\rho[I_t](\omega_{s'})}=\frac{\frac{1}{u_\rho'\left(e_\rho^{s}\right)}}{\frac{1}{u_\rho'\left(e_\rho^{s'}\right)}}\frac{\prod_{i=1}^{n}\left[u_i'\left(e_i^{s*}[I_t]+\theta_i^{s*}[I_t]\right)\pi_i[I_t](\omega_s)\right]^{\frac{\nu_i}{\sum_{i=1}^{n}\nu_i}}}{\prod_{i=1}^{n}\left[u_i'\left(e_i^{s'*}[I_t]+\theta_i^{s'*}[I_t]\right)\pi_i[I_t](\omega_{s'})\right]^{\frac{\nu_i}{\sum_{i=1}^{n}\nu_i}}}.$$

Box I.

Part (ii). Suppose now that there exists an equilibrium with price ratios given by (13). The equations

$$\frac{p_0^{s**}}{p_0^{s'**}}=\frac{p^{s**}[I_t]}{p^{s'**}[I_t]}$$

and

$$\frac{p^{s**}[I_t]}{p^{s'**}[I_t]}=\frac{u_i'\left(e_{i0}^{s}+\theta_{i0}^{s*}\right)\pi_i[I_t](\omega_s)}{u_i'\left(e_{i0}^{s'}+\theta_{i0}^{s'*}\right)\pi_i[I_t](\omega_{s'})}$$

hold for all $I_t\in\Pi_t$, $t\in\{0,\dots,T-1\}$ with $\omega_s,\omega_{s'}\in I_t$ if and only if

$$\frac{\pi_i[I_t](\omega_s)}{\pi_i[I_t](\omega_{s'})}=b$$

such that the constant $b>0$ is given as

$$b=\frac{u_i'\left(e_{i0}^{s'}+\theta_{i0}^{s'*}\right)}{u_i'\left(e_{i0}^{s}+\theta_{i0}^{s*}\right)}\frac{p^{s**}[I_t]}{p^{s'**}[I_t]}=\frac{u_i'\left(e_{i0}^{s'}+\theta_{i0}^{s'*}\right)}{u_i'\left(e_i^{s}[I_0]+\theta_i^{s*}[I_0]\right)}\frac{p_0^{s**}}{p_0^{s'**}}=\frac{\pi_i^0(\omega_s)}{\pi_i^0(\omega_{s'})}$$
$$=\frac{\frac{1}{\alpha[I_t]}\pi_i^0(\omega_s)}{\frac{1}{\alpha[I_t]}\pi_i^0(\omega_{s'})}$$

for any constant $\alpha[I_t]\neq 0$. Assume, to the contrary, that agent i is not a Bayesian decision maker, that is, assume that

$$\alpha[I_t]\neq\pi_i^0(I_t).$$

Then

$$\sum_{\omega\in I_t}\pi_i[I_t](\omega)=\frac{1}{\alpha[I_t]}\sum_{\omega\in I_t}\pi_i^0(\omega)=\frac{\pi_i^0(I_t)}{\alpha[I_t]}\neq 1,$$

which contradicts our assumption that $\pi_i[I_t]$ is a probability measure on $\left(I_t,2^{I_t}\right)$.□□

Proof of Theorem 1. Combining (12) with (15) establishes that there exists a representative agent model ρ for the naive equilibrium $\left(p_t^{**},\theta_t^{**}\right)_{t\in\{0,\dots,T\}}$ iff we have for all $I_t\in\Pi_t$, $t\in\{0,\dots,T-1\}$ and all $i\in\{1,\dots,n\}$

$$\frac{u_i'\left(e_i^{s*}[I_t]+\theta_i^{s*}[I_t]\right)\pi_i[I_t](\omega_s)}{u_i'\left(e_i^{s'*}[I_t]+\theta_i^{s'*}[I_t]\right)\pi_i[I_t](\omega_{s'})}=\frac{u_\rho'\left(e_\rho^{s}\right)}{u_\rho'\left(e_\rho^{s'}\right)}\frac{\pi_\rho[I_t](\omega_s)}{\pi_\rho[I_t](\omega_{s'})}\text{ whenever }\omega_s,\omega_{s'}\in I_t. \tag{27}$$

For arbitrary agent coefficients $\nu_i>0$ for $i\in\{1,\dots,n\}$, transform (27) equivalently to

$$\left(\frac{u_i'\left(e_i^{s*}[I_t]+\theta_i^{s*}[I_t]\right)\pi_i[I_t](\omega_s)}{u_i'\left(e_i^{s'*}[I_t]+\theta_i^{s'*}[I_t]\right)\pi_i[I_t](\omega_{s'})}\right)^{\nu_i}=\left(\frac{u_\rho'\left(e_\rho^{s}\right)}{u_\rho'\left(e_\rho^{s'}\right)}\frac{\pi_\rho[I_t](\omega_s)}{\pi_\rho[I_t](\omega_{s'})}\right)^{\nu_i}.$$

If this holds for all $i\in\{1,\dots,n\}$, we obtain (see Box I). Normalization, i.e., $\sum_{\omega_s\in I_t}\pi_\rho[I_t](\omega_s)=1$, yields the desired result.□□

Proof of Proposition 3. By Theorem 1, we are free to choose the agent coefficients $\nu_i>0$ in (16) arbitrarily. Fix the representative agent's CARA coefficient $\alpha_\rho>0$ and specify the agent coefficients as

$$\nu_i=\frac{\alpha_\rho}{\alpha_i}\text{ for all } i.$$

Observe that

$$\left(\prod_{i=1}^{n}\left(u'\left(e_i^{s*}[I_t]+\theta_i^{s*}[I_t]\right)\right)^{\nu_i}\right)^{\frac{1}{\sum_{i=1}^{n}\nu_i}}$$
$$=\left(\prod_{i=1}^{n}\alpha_i\exp\left(-\alpha_i\left(e_i^{s*}[I_t]+\theta_i^{s*}[I_t]\right)\frac{\alpha_\rho}{\alpha_i}\right)\right)^{\frac{1}{\sum_{i=1}^{n}\frac{\alpha_\rho}{\alpha_i}}}$$
$$=\left(\left(\prod_{i=1}^{n}\alpha_i\right)\exp\left(-\alpha_\rho\sum_{i=1}^{n}e_i^{s*}[I_t]+\theta_i^{s*}[I_t]\right)\right)^{\frac{1}{\sum_{i=1}^{n}\frac{\alpha_\rho}{\alpha_i}}}$$
$$=\left(\left(\prod_{i=1}^{n}\alpha_i\right)\exp\left(-\alpha_\rho\sum_{i=1}^{n}e_i^{s*}[I_t]+\theta_i^{s*}[I_t]\right)\right)^{\frac{1}{\sum_{i=1}^{n}\frac{\alpha_\rho}{\alpha_i}}}$$
$$=\left(\left(\prod_{i=1}^{n}\alpha_i\right)\exp\left(-\alpha_\rho\sum_{i=1}^{n}e_{i0}^{s}\right)\right)^{\frac{1}{\sum_{i=1}^{n}\frac{\alpha_\rho}{\alpha_i}}}$$

whereby the last step follows from

$$\sum_{i=1}^{n}e_{i0}^{s}=\sum_{i=1}^{n}e_i^{s*}[I_t]+\theta_i^{s*}[I_t]$$

whenever $\omega_s\in I_t$. For simplicity write $e_s=\sum_{i=1}^{n}e_{i0}^{s}$ for all s. Substituting in (16) yields (see Box II). Letting

$$e_\rho^{s}=e_s\frac{1}{\sum_{i=1}^{n}\frac{\alpha_\rho}{\alpha_i}}\text{ for all } s$$

gives us the desired result

$$\pi_\rho[I_t](\omega_s)=\frac{\left(\prod_{i=1}^{n}[\pi_i(\omega_s)]^{\frac{\alpha_\rho}{\alpha_i}}\right)^{\frac{1}{\sum_{i=1}\frac{\alpha_\rho}{\alpha_i}}}}{\sum_{s'=1}^{S}\left(\prod_{i=1}^{n}[\pi_i[I_t](\omega_{s'})]^{\frac{\alpha_\rho}{\alpha_i}}\right)^{\frac{1}{\sum_{i=1}\frac{\alpha_\rho}{\alpha_i}}}}.$$

□□

Proof of Proposition 4. Step 1. Observe that (11) can be equivalently rewritten for $i\in\{1,\dots,n\}$ as

$$\hat{\mu}_i[I_t]u'\left(e_i^{s*}[I_t]+\theta_i^{s*}[I_t]\right)\pi_i[I_t](\omega_s)=\hat{\mu}_i[I_t]\lambda_i[I_t]p^{s**}[I_t]$$
$$=\frac{1}{b}p^{s**}[I_t]$$

$$\pi_\rho\left[I_t\right](\omega_s) = \frac{\frac{1}{\exp(-\alpha_\rho e_{\rho s})}\left(\prod_{i=1}^n \alpha_i\right)^{\frac{1}{\sum_{i=1}^n \frac{\alpha_\rho}{\alpha_i}}}\left(\exp\left(-\alpha_\rho e_s\right)\right)^{\frac{1}{\sum_{i=1}^n \frac{\alpha_\rho}{\alpha_i}}}\left(\prod_{i=1}^n\left[\pi_i\left[I_t\right](\omega_s)\right]^{\frac{\alpha_\rho}{\alpha_i}}\right)^{\frac{1}{\sum_{i=1}\frac{\alpha_\rho}{\alpha_i}}}}{\sum_{s'=1}^S \frac{1}{\exp\left(-\alpha_\rho e_{\rho s'}\right)}\left(\prod_{i=1}^n \alpha_i\right)^{\frac{1}{\sum_{i=1}^n \frac{\alpha_\rho}{\alpha_i}}}\left(\exp\left(-\alpha_\rho e_{s'}\right)\right)^{\frac{1}{\sum_{i=1}^n \frac{\alpha_\rho}{\alpha_i}}}\left(\prod_{i=1}^n\left[\pi_i\left[I_t\right](\omega_{s'})\right]^{\frac{\alpha_\rho}{\alpha_i}}\right)^{\frac{1}{\sum_{i=1}\frac{\alpha_\rho}{\alpha_i}}}}$$

$$= \frac{\frac{1}{\exp(-\alpha_\rho e_{\rho s})}\exp\left(-\alpha_\rho e_s \frac{1}{\sum_{i=1}^n \frac{\alpha_\rho}{\alpha_i}}\right)\left(\prod_{i=1}^n\left[\pi_i(\omega_s)\right]^{\frac{\alpha_\rho}{\alpha_i}}\right)^{\frac{1}{\sum_{i=1}\frac{\alpha_\rho}{\alpha_i}}}}{\sum_{s'=1}^S \frac{1}{\exp\left(-\alpha_\rho e_{\rho s'}\right)}\exp\left(-\alpha_\rho e_{s'}\frac{1}{\sum_{i=1}^n \frac{\alpha_\rho}{\alpha_i}}\right)\left(\prod_{i=1}^n\left[\pi_i\left[I_t\right](\omega_{s'})\right]^{\frac{\alpha_\rho}{\alpha_i}}\right)^{\frac{1}{\sum_{i=1}\frac{\alpha_\rho}{\alpha_i}}}}.$$

Box II.

$$\pi_\rho\left[I_t\right](\omega_s) = \frac{\left(e_\rho^s\right)^\gamma \prod_{i=1}^n\left[\left(\frac{1}{a}\frac{\left(\hat{\mu}_i[I_t]\pi_i[I_t](\omega_s)\right)^{\frac{1}{\gamma}}}{\sum_{j=1}^n\left(\hat{\mu}_j[I_t]\pi_j[I_t](\omega_s)\right)^{\frac{1}{\gamma}}}e_\rho^s\right)^{-\gamma}\pi_i\left[I_t\right](\omega_s)\right]^{\frac{1}{n}}}{\sum_{\omega_{s'}\in I_t}\left(e_\rho^{s'}\right)^\gamma\prod_{i=1}^n\left[\left(\frac{1}{a}\frac{\left(\hat{\mu}_i[I_t]\pi_i[I_t](\omega_{s'})\right)^{\frac{1}{\gamma}}}{\sum_{j=1}^n\left(\hat{\mu}_j[I_t]\pi_j[I_t](\omega_{s'})\right)^{\frac{1}{\gamma}}}e_\rho^{s'}\right)^{-\gamma}\pi_i\left[I_t\right](\omega_{s'})\right]^{\frac{1}{n}}}$$

$$= \frac{\left(e_\rho^s\right)^\gamma\prod_{i=1}^n\left[\frac{\left(\hat{\mu}_i[I_t]\pi_i[I_t](\omega_s)\right)^{-1}}{\left(\sum_{j=1}^n\left(\hat{\mu}_j[I_t]\pi_j[I_t](\omega_s)\right)^{\frac{1}{\gamma}}\right)^{-\gamma}}\left(e_\rho^s\right)^{-\gamma}\right]^{\frac{1}{n}}\prod_{i=1}^n\left[\pi_i\left[I_t\right](\omega_s)\right]^{\frac{1}{n}}}{\sum_{\omega_{s'}\in I_t}\left(e_\rho^{s'}\right)^\gamma\prod_{i=1}^n\left[\frac{\left(\hat{\mu}_i[I_t]\pi_i[I_t](\omega_{s'})\right)^{-1}}{\left(\sum_{j=1}^n\left(\hat{\mu}_j[I_t]\pi_j[I_t](\omega_{s'})\right)^{\frac{1}{\gamma}}\right)^{-\gamma}}\left(e_\rho^{s'}\right)^{-\gamma}\right]^{\frac{1}{n}}\prod_{i=1}^n\left[\pi_i\left[I_t\right](\omega_{s'})\right]^{\frac{1}{n}}}$$

$$= \frac{\frac{\prod_{i=1}^n\left[\left(\hat{\mu}_i[I_t]\right)^{-1}\right]^{\frac{1}{n}}}{\left(\sum_{j=1}^n\left(\hat{\mu}_j[I_t]\pi_j[I_t](\omega_s)\right)^{\frac{1}{\gamma}}\right)^{-\gamma}}}{\sum_{\omega_{s'}\in I_t}\frac{\prod_{i=1}^n\left[\left(\hat{\mu}_i[I_t]\right)^{-1}\right]^{\frac{1}{n}}}{\left(\sum_{j=1}^n\left(\hat{\mu}_j[I_t]\pi_j[I_t](\omega_{s'})\right)^{\frac{1}{\gamma}}\right)^{-\gamma}}},$$

Box III.

where the agent weight $\hat{\mu}_i\left[I_t\right]$ is given as (20). Using

$$\frac{\hat{\mu}_i\left[I_t\right]u'\left(e_i^{s*}\left[I_t\right]+\theta_i^{s*}\left[I_t\right]\right)\pi_i\left[I_t\right](\omega_s)}{\hat{\mu}_j\left[I_t\right]u'\left(e_j^{s*}\left[I_t\right]+\theta_j^{s*}\left[I_t\right]\right)\pi_j\left[I_t\right](\omega_s)} = \frac{\frac{1}{b}p^{s**}\left[I_t\right]}{\frac{1}{b}p^{s**}\left[I_t\right]} = 1$$

yields for CRRA Bernoulli utility

$$\frac{\hat{\mu}_i\left[I_t\right]}{\hat{\mu}_j\left[I_t\right]} = \left(\frac{e_i^{s*}\left[I_t\right]+\theta_i^{s*}\left[I_t\right]}{e_j^{s*}\left[I_t\right]+\theta_j^{s*}\left[I_t\right]}\right)^\gamma\frac{\pi_j\left[I_t\right](\omega_s)}{\pi_i\left[I_t\right](\omega_s)}$$

$$\Leftrightarrow$$

$$\left(\frac{\hat{\mu}_i\left[I_t\right]\pi_i\left[I_t\right](\omega_s)}{\hat{\mu}_j\left[I_t\right]\pi_j\left[I_t\right](\omega_s)}\right)^{\frac{1}{\gamma}} = \frac{e_i^{s*}\left[I_t\right]+\theta_i^{s*}\left[I_t\right]}{e_j^{s*}\left[I_t\right]+\theta_j^{s*}\left[I_t\right]}.$$

Let $e_\rho^s = ae_s = a\sum_{i=1}^n\left(e_i^{s*}\left[I_t\right]+\theta_i^{s*}\left[I_t\right]\right)$ for an arbitrary $a > 0$. Then this system of equations has the solution

$$e_i^{s*}\left[I_t\right]+\theta_i^{s*}\left[I_t\right] = \frac{1}{a}\frac{\left(\hat{\mu}_i\left[I_t\right]\pi_i\left[I_t\right](\omega_s)\right)^{\frac{1}{\gamma}}}{\sum_{j=1}^n\left(\hat{\mu}_j\left[I_t\right]\pi_j\left[I_t\right](\omega_s)\right)^{\frac{1}{\gamma}}}e_\rho^s. \tag{28}$$

Step 2. By Theorem 1, we are free to choose arbitrary agent coefficients in (16). Simply let $\nu_i = 1$ for all i. Substituting (28) in the general belief aggregation formula (16) gives us (see Box III)

which yields the desired result

$$\pi_\rho\left[I_t\right](\omega_s) = \frac{\left(\sum_{j=1}^n\left(\hat{\mu}_j\left[I_t\right]\pi_j\left[I_t\right](\omega_s)\right)^{\frac{1}{\gamma}}\right)^\gamma}{\sum_{\omega_{s'}\in I_t}\left(\sum_{j=1}^n\left(\hat{\mu}_j\left[I_t\right]\pi_j\left[I_t\right](\omega_{s'})\right)^{\frac{1}{\gamma}}\right)^\gamma}.$$

□□

Proof of Proposition 5. Step 1. Given the state-price vector p_0, define the probability measure Q on $\left(\Omega, 2^\Omega\right)$ such that, for all $\omega_s \in \Omega$,

$$Q\left(\omega_s\right) = \frac{p_0^s}{\sum_{\{s'|\omega_{s'}\in\Omega\}}p_0^{s'}} > 0.$$

By construction, we have, for all $\omega_s \in \Omega$, the equivalence

$$\frac{p_0^s}{p_0^{s'}} = \frac{Q\left(\omega_s\right)}{Q\left(\omega_{s'}\right)} = \frac{\sum_{\omega\in\Omega}\mathbf{1}_{\{\omega_s\}}(\omega)\,Q(\omega)}{\sum_{\omega\in\Omega}\mathbf{1}_{\{\omega_{s'}\}}(\omega)\,Q(\omega)}$$

$$= \frac{\mathbb{E}_Q p^s[I_T]}{\mathbb{E}_Q p^{s'}[I_T]}.$$

Step 2. Fix an arbitrary $t < T$. Let $\omega_s, \omega_{s'} \in I_t$ so that (24) is, by Step 1, equivalent to

$$\frac{p_t^s[I_t]}{p_t^{s'}[I_t]} = \frac{\mathbb{E}_Q p_T^s}{\mathbb{E}_Q p_T^{s'}}$$
$$= \frac{\sum_{I_t\in\Pi_t}\left(\sum_{\omega\in\Omega}\mathbf{1}_{\{\omega_s\}}(\omega)\,Q(\omega\mid I_t)\right)Q(I_t)}{\sum_{I_t\in\Pi_t}\left(\sum_{\omega\in\Omega}\mathbf{1}_{\{\omega_{s'}\}}(\omega)\,Q(\omega\mid I_t)\right)Q(I_t)}$$
by the law of iterated expectations
$$= \frac{\sum_{\omega\in\Omega}\mathbf{1}_{\{\omega_s\}}(\omega)\,Q(\omega\mid I_t)\,Q(I_t)}{\sum_{\omega\in\Omega}\mathbf{1}_{\{\omega_{s'}\}}(\omega)\,Q(\omega\mid I_t)\,Q(I_t)}$$
$$= \frac{\mathbb{E}_{Q(\cdot|I_t)}p^s[I_T]}{\mathbb{E}_{Q(\cdot|I_t)}p^{s'}[I_T]}.$$

Denote by α_t an arbitrary $\sigma(\Pi_t)$ -measurable random variable such that $\alpha_t[I_t] > 0$ for all $I_t \in \Pi_t$. Set the price

$$p^s[I_t] = \alpha_t[I_t]\,\mathbb{E}_{Q(\cdot|I_t)}p^s[I_T] \quad (29)$$

and note that

$$\mathbb{E}_{Q(\cdot|I_t)}p^s[I_T] > 0 \text{ iff } \omega_s \in I_t$$

so that

$$\frac{p^s[I_t]}{p^{s'}[I_t]} = \frac{\mathbb{E}_{Q(\cdot|I_t)}p^s[I_T]}{\mathbb{E}_{Q(\cdot|I_t)}p^{s'}[I_T]}$$

holds whenever $\omega_s, \omega_{s'} \in I_t$ whereby $p^s[I_t] = 0$ iff $\omega_s \notin I_t$.

Step 3. Applying the pricing rule (29) to $t+1$ yields

$$p^s[I_{t+1}] = \alpha_{t+1}[I_{t+1}]\,\mathbb{E}_{Q(\cdot|I_{t+1})}p^s[I_T].$$

If (and only if) the random variable

$$\frac{\alpha_{t+1}[I_{t+1}]}{\alpha_t[I_t]}$$

is $\sigma(\Pi_t)$-measurable (i.e., taking on a constant value on I_t), the conditional expectation of $p^s[I_{t+1}]$ becomes

$$\mathbb{E}_{Q(\cdot|I_t)}p^s[I_{t+1}] = \sum_{I_{t+1}\subseteq I_t} p^s[I_{t+1}]\,Q(I_{t+1}\mid I_t)$$
$$= \sum_{I_{t+1}\subseteq I_t}\alpha_{t+1}[I_{t+1}]\,\mathbb{E}_{Q(\cdot|I_{t+1})}p^s[I_T]\,Q(I_{t+1}\mid I_t)$$
$$= \sum_{I_{t+1}\subseteq I_t}\frac{\alpha_{t+1}[I_{t+1}]}{\alpha_t[I_t]}\alpha_t[I_t]\,\mathbb{E}_{Q(\cdot|I_{t+1})}p^s[I_T]\,Q(I_{t+1}\mid I_t)$$
$$= \frac{\alpha_{t+1}[I_{t+1}]}{\alpha_t[I_t]}\alpha_t[I_t]\sum_{I_{t+1}\subseteq I_t}\mathbb{E}_{Q(\cdot|I_{t+1})}p^s[I_T]\,Q(I_{t+1}\mid I_t)$$
$$= \frac{\alpha_{t+1}[I_{t+1}]}{\alpha_t[I_t]}\alpha_t[I_t]\,\mathbb{E}_{Q(\cdot|I_t)}p^s[I_T]$$
by the law of iterated expectations
$$= \frac{\alpha_{t+1}[I_{t+1}]}{\alpha_t[I_t]}p^s[I_t],$$

whereby the last step follows from (29).

Step 4. By Step 3, we have constructed an $(\mathcal{F}_t)_{t\in\{0,\ldots,T\}}$-adapted price process $\left(p_t^s\right)_{t\in\{0,\ldots,T\}}$ such that for all $I_t \in \Pi_t$, $t \in \{0,\ldots,T-1\}$,

$$p^s[I_t] = \frac{\alpha_t[I_t]}{\alpha_{t+1}[I_{t+1}]}\mathbb{E}_{Q(\cdot|I_t)}p^s[I_{t+1}] \quad (30)$$

where $\frac{\alpha_t[I_t]}{\alpha_{t+1}[I_{t+1}]}$ is an arbitrary strictly positive, $\sigma(\Pi_t)$-measurable random variable. Define for all $I_t \in \Pi_t$, $t \in \{0,\ldots,T-1\}$

$$r_{t,t+1}[I_t] = \frac{\alpha_{t+1}[I_{t+1}]}{\alpha_t[I_t]} - 1$$

as well as $r_{T,T+1}[I_T(\omega)] = 0$ for all $\omega \in \Omega$ to obtain the desired (discounted) martingale

$$p^s[I_t] = \frac{1}{1 + r_{t,t+1}[I_t]}\mathbb{E}_{Q(\cdot|I_t)}p^s_{t+1}.$$

where $\left(r_{t,t+1}\right)_{t\in\{0,\ldots,T\}}$ is an arbitrary $(\mathcal{F}_t)_{t\in\{0,\ldots,T\}}$ -adapted short-rate process with $1 + r_{t,t+1}[I_t(\omega)] > 0$ for all $\omega \in \Omega$ for $t < T$.□□

References

Arrow, K.J., 1974. Essays in the Theory of Risk-Bearing. North-Holland Publishing Company, Amsterdam.
Atmaz, A., Basak, S., 2018. Belief dispersion in the stock market. J. Finance 73, 1225–1279.
Baker, R., 2022. Mathematical models of confirmation bias working paper. p. 2022, arXiv preprint arXiv:2202.03072.
Barberis, N., Shleifer, A., Vishny, R., 1998. A model of investor sentiment. J. Financ. Econ. 49, 307–343.
Berk, R.H., 1966. Limiting behaviour of posterior distributions when the model is incorrect. Ann. Math. Stat. 31, 51–58.
Bhamra, H.S., Uppal, R., 2014. Asset prices with heterogeneity in preferences and beliefs. Rev. Financ. Stud. 27, 519–580.
Billingsley, P., 1995. Probability and Measure. Wiley, New York.
Blume, L., Easley, D., 2006. If you're so smart, why aren't you rich? Belief selection in complete and incomplete markets. Econometrica 74, 929–966.
Calvet, L.-E., Grandmont, J.-M., Lemaire, I., 2018. Aggregation of heterogenous beliefs, asset pricing, and risk sharing in complete financial markets. Res. Econ. 72, 117–146.
Dana, R.-A., 1993a. Existence and uniqueness of equilibria when preferences are additively separable. Econometrica 61, 953–957.
Dana, R.-A., 1993b. Existence, uniqueness and determinacy of Arrow–Debreu equilibria in finance models. J. Math. Econom. 22, 563–579.
Daniel, K.D., Hirshleifer, D., Subrahmanyam, A., 1998. Investor psychology and security market under- and overreactions. J. Finance 53, 1839–1885.
Doob, J.L., 1949. Application of the theory of martingales. In: Le Calcul des Probabilités et ses Applications, Vol. 13. Colloques Internationaux du Centre National de la Recherche Scientifique, pp. 23–27.
Duffie, D., 2001. Dynamic Asset Pricing Theory. Princeton University Press, Princeton.
Epstein, L.G., 2006. An axiomatic model of non-Bayesian updating. Rev. Econom. Stud. 73, 413–436.
Epstein, L.G., Le Breton, M., 1993. Dynamically consistent beliefs must be Bayesian. J. Econom. Theory 61, 1–22.
Epstein, L.G., Noor, J., Sandroni, A., 2008. Non-Bayesian updating: A theoretical framework. Theor. Econ. 3, 193–229.
Epstein, L.G., Schneider, M., 2003. Recursive multiple-priors. J. Econom. Theory 113, 1–31.
Gennaioli, N., Shleifer, A., 2010. What comes to mind. Q. J. Econ. 125, 1399–1433.
Ghirardato, P., 2002. Revisiting Savage in a conditional world. Econom. Theory 20, 83–92.
Gilboa, I., Schmeidler, D., 1993. Updating ambiguous beliefs. J. Econom. Theory 59, 33–49.
Groneck, M., Ludwig, A., Zimper, A., 2022. Who saves more, the naive or the sophisticated agent? In: Working Paper. Downloadable at: https://alexander-ludwig.com/wp-content/uploads/2022/07/NaiveSoph_Main_Long_012022.pdf.
Guerdjikova, A., Quiggin, J., 2019. Market selection with differential financial constraints. Econometrica 87, 1693–1762.
Harrison, J.M., Kreps, D.M., 1979. Martingales and arbitrage in multiperiod securities markets. J. Econom. Theory 20, 381–408.
Hong, H., Stein, J.C., 1999. A unified theory of underreaction, momentum trading, and overreaction in asset markets. J. Finance 54, 2143–2184.
Jouini, E., Napp, C., 2007. Consensus consumer and intertemporal asset pricing with heterogeneous beliefs. Rev. Econom. Stud. 74, 1149–1174.
Kullback, S., Leibler, R.A., 1951. On information and sufficiency. Ann. Math. Stat. 22, 79–86.
Lucas, Jr., R.E., 1978. Asset prices in an exchange economy. Econometrica 46, 1429–1445.
Ludwig, A., Zimper, A., 2013. A decision-theoretic model of asset-price underreaction and overreaction to dividend news. Ann. Finance 9, 625–665.
Martin, I.W.R., Papadimitriou, D., 2021. Sentiment and speculation in a market with heterogeneous beliefs. In: Working Paper.
Mas-Colell, A., 1985. TMarket with Heterogenral Economic Equilibrium: A Differentiable Approach. In: Econometric Society Monographs, Cambridge University Press, Cambridge.

Mas-Colell, A., Whinston, M.D., Green, J.R., 1995. Microeconomic Theory. Oxford University Press, Oxford.
Mullainathan, S., Schwartzstein, J., Shleifer, A., 2008. Coarse thinking and persuasion. Q. J. Econ. 123, 577–619.
Negishi, T., 1960. Welfare economics and existence of an equilibrium for a competitive economy. Metroeconomica 12, 92–97.
O'Donoghue, T., Rabin, M., 1999. Doing it now or later. Am. Econ. Rev. 89, 103–124.
Ortoleva, P., 2012. Modeling the change of paradigm: Non-Bayesian reactions to unexpected news. Am. Econ. Rev. 102, 2410–2436.
Rabin, M., 2002. Inference by believers in the law of small numbers. Q. J. Econ. 117, 775–816.
Rabin, M., Schrag, J.L., 1999. First impressions matter: A model of confirmatory bias. Q. J. Econ. 114, 37–82.
Rubinstein, M., 1974. An aggregation theorem for securities markets. J. Financ. Econ. 1, 225–244.
Rubinstein, M., 1976. The strong case for the generalized logarithmic utility model as the premier model of financial markets. J. Finance 31, 551–571.
Sandroni, A., 2000. Do markets favor agents able to make accurate predictions? Econometrica 68, 1303–1342.
Savage, L.J., 1954. The Foundations of Statistics. Wiley, New York.
Siniscalchi, M., 2011. Dynamic choice under ambiguity. Theor. Econ. 6, 379–421.
Zimper, A., 2023. Belief aggregation for representative agent models. Int. J. Econ. Theory 19, 309–342.
Zimper, A., Ma, W., 2017. Bayesian learning with multiple priors and non-vanishing ambiguity. Econom. Theory 64, 409–447.