Belief-neutral efficiency in financial markets*

Patrick Beißner

Research School of Economics
The Australian National University
patrick.beissner@anu.edu.au

Frank Riedel

Center for Mathematical Economics
Bielefeld University
frank.riedel@uni-bielefeld.de
and
School of Economics
University of Johannesburg

February 2025

Abstract

Heterogeneous beliefs among market participants can lead to questionable speculative trading that goes beyond any risk-sharing motives. We demonstrate that such unwarranted betting behavior in market equilibrium can be mitigated by introducing nonlinear pricing for ambiguous contracts, without compromising legitimate risk-hedging activities. While Arrow-Debreu equilibria generally fail to achieve belief-neutral efficiency, we establish a modified version of the first welfare theorem in which equilibria with nonlinear prices uphold belief-neutral efficiency. Moreover, we show that belief-neutral efficiency can be ensured by introducing suitable transaction costs for ambiguous financial assets.

Key words and phrases: Belief-Neutral Pareto efficiency, Heterogeneous Beliefs, Knightian Uncertainty, Ambiguity, General Equilibrium

JEL subject classification: D81, C61, G11

*The authors gratefully acknowledge the financial support of the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – SFB 1283/2 2021 – 317210226 (Project C5).

1 Introduction

Excessive trading in financial markets is a well-documented phenomenon (Odean (1999)) and poses a puzzle when compared to the minimal trading required to achieve intertemporal efficiency in stationary environments (Judd, Kubler, and Schmedders (2003)).

We develop a stylized model where excessive trading arises from heterogeneous beliefs under Knightian uncertainty. While all economic fundamentals have commonly known probabilities, placing the economy in a setting of risk, there is additional uncertainty about unrelated events, such as sports outcomes or the success of new startups.

In this economy, endowments are risky but not uncertain, as agents share common beliefs about fundamental probabilities. However, they encounter external betting opportunities where their beliefs diverge, driven by Knightian uncertainty or behavioral biases like overconfidence.

Agents may engage in unnecessary betting in such cases, and we demonstrate that this behavior generically occurs in Arrow-Debreu equilibrium. As a result, these equilibria are not belief-neutrally efficient in the sense of Brunnermeier, Simsek, and Xiong (2014).

This fact raises the question of whether appropriate market design can lead to better outcomes. We show that incorporating nonlinear pricing, as proposed by Beissner and Riedel (2019), which accounts for the ambiguity of betting events, eliminates undesirable betting. The resulting *Knight–Walras* equilibria are belief-neutrally efficient, while prices for fundamentals linked to individuals' idiosyncratic risks remain linear, ensuring the desired risk-sharing trade.

As the notion of a *Knight-Walras* equilibrium is somewhat abstract, we ask which kind of rules might implement the desired equilibrium in practice. We show that suitable designed transaction costs for ambiguous assets and events do the job. The resulting financial markets remains frictionless for purely hedging related trade and thus does not hamper social welfare of reasonable, not betting-related trade.

Literature. There is a large and ongoing literature on heterogeneous beliefs in economics and finance that can emerge from various sources: Investors can have heterogeneous prior beliefs (Morris (1995)), they might interpret information heterogeneously (Harris and Raviv (1993), Kandel and Pearson (1995), Banerjee and Kremer (2010)). The literature on belief formation has increasingly emphasized the role of behavioral forces in shaping heterogeneously distorted beliefs among individuals. Overconfidence, for instance, has been extensively studied in financial markets, where individuals tend to overestimate their abilities or the precision of their information, leading to suboptimal decisions (Odean (1998), Daniel, Hirshleifer, and Subrahmanyam

(2001)). Similarly, motivated reasoning, wherein individuals selectively interpret information to reinforce their pre-existing beliefs or desires, has been shown to distort decision-making in economic and political contexts (Brunnermeier and Parker (2005), Bénabou and Tirole (2016)). This paper builds on these insights to further explore the implications of heterogeneous belief distortions for economic outcomes.

Economists have long recognized that the classic Pareto criterion often leads to unappealing outcomes when agents hold heterogeneous beliefs. Early general equilibrium literature include von Weizsäcker (1969), Dreze (1972), Starr (1973), Harris (1978), Hammond (1981). More recently – in the aftermath of the great financial crisis – a new interest in efficiency criteria under heterogeneous beliefs emerged, see Gilboa, Samuelson, and Schmeidler (2014), Brunnermeier, Simsek, and Xiong (2014), Danan, Gajdos, Hill, and Tallon (2016).

A consequence of heterogeneous belief in markets often leads to excessive trade and speculation. In the setting of financial market equilibria this is well studied, see Harrison and Kreps (1978), Detemple and Murthy (1994), Scheinkman and Xiong (2003), Jouini and Napp (2007), and Borovička (2020). The rationale for introducing a new efficiency notion relies on the desire to understand a normative criterion that mitigates such socially nondesirable equilibrium outcome.

As highlighted by Stiglitz (1989), the welfare-diminishing effect of speculation in financial markets may warrant a corrective market design. Rather than commonly proposed transaction taxes, our results—assuming belief-neutral efficiency as the normative benchmark—suggest an alternative approach: a specific form of market incompleteness regulated through bid-ask spreads. In the same vein, Blume, Cogley, Easley, Sargent, and Tsyrennikov (2018) advocate how particular forms of market incompleteness can potentially mute some forms of excessive trade, see also Simsek (2013) for an inverse approach. Towards this end our suggested design approach also allows for a link to the literature on market regulation, see Athanasoulis and Shiller (2000), Easley and O'Hara (2009), Posner and Weyl (2013).

Our paper is also related to the huge historical debate about appropriate transaction costs for financial markets, dating back at least to Tobin. In a sense, we provide a general equilibrium foundation for such transaction cost; in a related vein, Dávila (2023) has recently studied the optimal transaction tax in avery specific mean-variance setting.

The paper is organized as follows. Section 2 presents the basic economy. Section 3 introduces the belief-neutral efficiency notion and its inefficiency in Arrow-Debreu equilibira. Section 4 establish belief-neutral efficiency in Knight-Walras equilibria. Section 5 establishes what financial market allow for belief-neutral efficiency. The final section concludes. The Appendix contains the proofs and details of our leading example.

2 The Economy

We examine the difference of purely speculative and insurance-oriented trading within a framework that differentiates between fundamental risks—where agents share common beliefs—and subjective uncertainty, where beliefs may diverge.

There are finitely many states of the world S; we denote the power set of S by $\mathcal{F}$. The commodity space is given by $\mathbb{X} = \mathbb{R}^S$ and the consumption set by its positive cone $\mathbb{X}_+ = \mathbb{R}^S_+$. Denote by ΔS the probability simplex on S.

The economy consists of finitely many agents $i = 1, \ldots, I$ with subjective beliefs $P_i \in \Delta S$. Agents' risk aversion is described by a Bernoulli utility function $u_i : \mathbb{R}_+ \to \mathbb{R}$. Their preferences are thus represented by the subjective expected utility function[1]

$$U_i(c) = E^{P_i}[u_i(c)]$$

for $c \in \mathbb{X}_+$. The Bernoulli utility functions satisfy standard assumptions as, e.g., in Dana (1993): u_i is strictly increasing and strictly concave, twice continuously differentiable on $(0, \infty)$, and $\lim_{c \downarrow 0} u_i'(c) = \infty$. We assume throughout that the individual endowments $e_i : S \to \mathbb{R}_+$ are strictly positive in every state of the world.

We are thus in a situation of Knightian uncertainty in as far as agents' beliefs disagree on some events of the world. We denote by $\mathcal{P}$ the convex hull of the individual beliefs P_i. We denote by $\mathcal{G}$ the algebra of events on which agents agree, i.e. we have $A \in \mathcal{G}$ if and only if for all i and j we have $P_i(A) = P_j(A)$. We call $\mathcal{G}$ the algebra of risky events[2] and we assume that $\mathcal{G}$ is strictly smaller than the power set $\mathcal{F}$ of S.

We aim to study a situation in which agents agree on the fundamental risks of the economy but may disagree about the probability of some events that do not affect the fundamental economic activities. To this end, we assume that there is no Knightian uncertainty about the relevant fundamentals of our economy; the initial endowments $e_i \in \mathbb{X}_+$, $\mathcal{G}$–measurable. There is thus no disagreement about the probability distribution of endowments.

Keeping the other ingredients fixed, we write $\mathcal{E}(P_1, \ldots, P_I)$ for the economy with belief profile $(P_1, \ldots, P_I)$.

[1] We conjecture that our main results hold true for ambiguity-averse preferences as long as the subjective beliefs in the sense of Rigotti, Shannon, and Strzalecki (2008) are sufficiently diverse. The point we make is easiest to see with ambiguity-neutral agents.

[2] The algebra appears under different names in the literature. For example, Ghirardato, Maccheroni, and Marinacci (2004) call such events *crisp*. Gul and Pesendorfer (2014) speak of *ideal* events.

3 Belief-Neutral Inefficiency of Markets

In this section, we address the common occurrence of unjustified betting under subjective heterogeneous beliefs. We begin by reviewing the standard concepts of efficiency and equilibrium. Let $e = \sum_i e_i$ be the aggregate endowment;

$$\Lambda(e) = \left\{ (x_i) \in \mathbb{X}_+^I : \sum_i x_i \leq e \right\}$$

is the set of feasible allocations. A feasible allocation $(c_i) \in \Lambda(e)$ is efficient if there is no other feasible allocation $(y_i) \in \Lambda(e)$ that makes every agent better off in the sense[3] that $U_i(y_i) > U_i(c_i)$. An Arrow–Debreu equilibrium consists of a linear price functional $\Psi : \mathbb{X} \to \mathbb{R}$ and a feasible allocation $(c_i) \in \Lambda(e)$ such that for all agents i and budget-feasible consumption plans y with $\Psi(y - e_i) \leq 0$ we have $U_i(y) \leq U_i(c_i)$.

With heterogeneous beliefs, an alternative notion of efficiency has been proposed by Brunnermeier, Simsek, and Xiong (2014). The social planner recognizes the presence of ambiguity (or the diversity of potentially distorted beliefs), and acknowledges that all of them belong to the class of reasonable beliefs. Eficiency is required to be robust across all of the reasonable beliefs. A feasible allocation $(c_i) \in \Lambda(e)$ is called *belief-neutral efficient* if it is efficient in every homogeneous belief economy $\mathcal{E}(Q, \ldots, Q)$ for all reasonable beliefs $Q \in \mathcal{P}$.

We next illustrate how unjustified speculative trade can arise in an Arrow-Debreu equilibrium with the help of a simple example that illustrates nicely the basic problem.

Example 1 *To make things explicit, assume four states*

$$S = \{s_1, s_2, s_A, s_B\}$$

and consider two agents. States s_1 and s_2 are the risky states where agents' beliefs agree, e.g. we have $P_i(s_1) = P_i(s_2) = 1/4$ for agents $i = 1, 2$. However, agent 1 thinks that the state s_A, maybe related to a sports event in which country A beats country B, has a higher probability than state s_B, say $P_1(s_A) = 3/8$, whereas we have exactly the opposite for agent 2, so $P_2(s_B) = 3/8$. Consider constant relative risk aversion of 1 for all agents, i.e. $u_i(x) = \log(x)$. The economy's endowments vary only in the first two states, say

$$e_1(s_1) = 3, e_1(s_2) = 1, e_1(s_A) = e_1(s_B) = 2$$

[3]With our assumptions on Bernoulli utility functions, one can show that the weak form of efficiency that we use is equivalent to the usual one in which one requires the strict inequality only for one agent. See the appendix for details.

and

$$e_2(s_1) = 1, e_1(s_2) = 3, e_1(s_A) = e_1(s_B) = 2.$$

The symmetry of the situation suggests that a good allocation would be the full insurance allocation in which both agents obtain two units of the consumption good in each state. This allocation would indeed be the equilibrium allocation if agents had uniform beliefs. Let us consider what happens in Arrow-Debreu equilibrium with distinct beliefs about the sports' related states s_A and s_B.

Write the price functional in the form

$$\Psi(x) = \sum_{s \in S} \psi(s)x(s)$$

for $x \in \mathbb{X}$. The first-order condition is

$$P_i(s)\frac{1}{c_i(s)} = \lambda_i \psi(s)$$

for a Lagrange multiplier λ_i. The budget constraint yields $\lambda_i = 1/\Psi(e_i)$. Hence, we obtain

$$c_i(s) = \Psi(e_i)\frac{P_i(s)}{\psi(s)}.$$

In equilibrium, markets clear, so

$$4 = \sum_{i=1}^{2} c_i(s) = \frac{1}{\psi(s)} \sum_{i=1}^{2} \Psi(e_i)P_i(s),$$

resulting in the equilibrium state price

$$\psi(s) = \frac{1}{4} \sum_{i=1}^{2} \Psi(e_i)P_i(s).$$

We thus see that the equilibrium state price is a convex combination of agents' beliefs. In fact, if we take $\psi(s) = 1/4$ for all states, we have found the equilibrium price, as one easily checks, with symmetric wealth $\Psi(e_i) = 2$. This solution is also the unique solution of the equilibrium conditions.

As a consequence, in the risky states s_1 and s_2, both agents fully insure and consume two units in each state, whereas in the speculative states, where beliefs differ, they start speculating, with agents consuming

$$c_1(s_A) = \Psi(e_i)\frac{P_1(s_A)}{\psi(s_A)} = 3, c_2(s_A) = 1, c_1(s_B) = 1, c_2(s_B) = 3.$$

Agents have insured the risk of the first two states, but have created new, unnecessary risks on the sports events.

The Arrow-Debreu allocation is not belief-neutral efficient because for any reasonable belief of the form $P^h = \sum_{i=1}^{2} h_i P_i$, the allocation c is dominated by the full insurance allocation

$$\bar{c}_i^h = E^{P^h} c_i$$

as one can easily prove with the help of Jensen's inequality (or direct computation, for that matter).

Although the above example is very special, nothing is special about unjustified trade occurring in markets. Indeed, belief-neutral inefficient trade occurs generically in Arrow-Debreu equilibria when agents hold heterogeneous beliefs as the next theorem shows.

Theorem 1 *Generically in beliefs, Arrow-Debreu equilibrium allocations are belief-neutral inefficient.*

A rigorous formulation of the statement in measure-theoretic terms as well as the proof is found in Appendix A.

The basic intuition why belief-neutral efficiency fails in markets is as follows. Belief-neutral efficient allocation equate agents' marginal rates of substitution state by state without referring to beliefs,

$$\frac{u_i'\left(c_i^*(s)\right)}{u_i'\left(c_i^*(t)\right)} = \frac{u_j'\left(c_j^*(s)\right)}{u_j'\left(c_j^*(t)\right)}$$

for all agents i, j and states s, t. With heterogeneous beliefs, however, the first-order conditions imply that in equilibrium we have instead

$$\frac{P_i(s)u_i'\left(c_i^*(s)\right)}{P_i(t)u_i'\left(c_i^*(t)\right)} = \frac{P_j(s)u_j'\left(c_j^*(s)\right)}{P_j(t)u_j'\left(c_j^*(t)\right)} \tag{1}$$

for all agents i, j and states s, t. With distinct beliefs, the two systems of equations are not consistent with each other.

We now present a result on belief-neutral efficiency that might be of independent interest. Let us first recall the characterization of interior efficient allocations in homogeneous belief economies. For individual weights $\alpha \in \Delta I$, we define the representative agent's Bernoulli utility of the homogeneous belief economy

$$u_\alpha(x) = \max_{c_i \in \mathbb{R}_+^I : \sum_i c_i = x} \sum_i \alpha_i u_i(c_i) \tag{2}$$

for $x > 0$. Under our standing assumptions, u_α is continuously differentiable. The unique maximizers in (2) are given by continuously differentiable functions $C_{\alpha,i}(x)$ that satisfy the first-order conditions

$$\alpha_i u_i'\left(C_{\alpha,i}(x)\right) = u_\alpha'(x). \tag{3}$$

These properties are well known, see, e.g. Dana (1993). We now characterize belief-neutral efficient allocations.

Lemma 1 *Let $(c_i^*) \in \Lambda(e)$ be an interior feasible allocation. The following assertions are equivalent*

1. *(c_i^*) is belief neutral efficient.*
2. *The allocation satisfies*

$$c_i^*(s) = C_{\alpha,i}(e(s))$$

for some weights $\alpha_i > 0$ and continuously differentiable functions $C_{\alpha,i} : (0,\infty) \to \mathbb{R}_+$ that satisfy (3).

Brunnermeier, Simsek, and Xiong (2014) characterize belief-neutral efficiency in Proposition 1 with the help of Bergson-Samuelson social welfare functions. The above Lemma 1 yields an analytical tractable formulation.

4 Belief-neutral efficient equilibria

In economies characterized by belief heterogeneity, the competitive market outcome almost never achieves belief-neutral efficiency. This raises a natural question: can belief-neutral efficiency be attained through appropriate market mechanisms? In this section, we show that implementing sublinear pricing on events where agents hold divergent beliefs can effectively induce belief-neutral efficiency.

Consider the sublinear expectation

$$\mathbb{E}^{\mathcal{P}} x = \max_{P \in \mathcal{P}} \mathbb{E}^P x.$$

This sublinear expectation is linear on the set of all $\mathcal{G}$-measurable consumption plans because all $P \in \mathcal{P}$ coincide on $\mathcal{G}$-measurable events. It is nonlinear outside this set, in general.

In previous work (Beissner and Riedel (2019)), we defined an equilibrium concept based on sublinear expectations of the form

$$\mathbb{E}^{\mathcal{Q}} x = \max_{P \in \mathcal{Q}} \mathbb{E}^P x$$

for $x \in \mathbb{X}$ and a closed convex set of probabilities $\mathcal{Q} \subset \Delta S$ and we developed the corresponding equilibrium theory. A pair (ψ, c) consisting of a state–price $\psi : \Omega \to \mathbb{R}_+$ and a feasible allocation $c = (c_i)_{i=1,\dots,I} \in \Lambda(e)$ is a Knight-Walras equilibrium (with respect to $\mathcal{Q}$) if for each i, c_i maximizes agent i's utility subject to the Knight-Walras budget constraint

$$\Psi^{\mathcal{Q}}(c-e_i) := \max_{Q\in\mathcal{Q}} E^Q[\psi(c-e_i)] \leq 0. \tag{4}$$

We now come to our main theorem that shows how introducing nonlinear prices on payoffs where agents' beliefs diverge prevents unjustified betting.

Theorem 2 *In every Knight-Walras equilibrium (ψ^*, c^*) with respect to the set of reasonable beliefs $\mathcal{Q} = \mathcal{P}$, the equilibrium allocation c^* is belief-neutral efficient.*

The proof can be found in the appendix. To see how nonlinear prices on ambiguous trades work, let us go back to our running example.

Example 1 *(continued) The convex hull $\mathcal{P}$ of the two beliefs P_1 and P_2 can be parametrized as follows, see also Figure 1,*

$$\mathcal{P} = \left\{ \left(\tfrac{1}{4}, \tfrac{1}{4}, p, \tfrac{1}{2} - p\right) : p \in \left[\tfrac{1}{8}, \tfrac{3}{8}\right] \right\}.$$

In Knight-Walras equilibrium, we have $\psi(s) = 1$ in all states and agents fully insure, consuming $c_i^*(s) = 2$ in all states, see Appendix D for details of the computation.

The Knight-Walras equilibrium price system is given by

$$\begin{aligned}\Psi^{\mathcal{P}}(x_1, x_2, x_A, x_B) &= \max_{p\in\left[\frac{1}{8},\frac{3}{8}\right]} \left(\frac{x_1 + x_2}{4} + p x_A + \left(\frac{1}{2} - p\right) x_B \right) \\ &= \frac{x_1 + x_2 + x_A + x_B}{4} + \frac{1}{8}(x_A - x_B)^+ + \frac{1}{8}(x_B - x_A)^+ .\end{aligned}$$

The price system is linear in x_1 and x_2, in consumption in the two risky states where beliefs agree, and nonlinear in consumption in x_A and x_B where beliefs diverge.

The speculative Arrow-Debreu equilibrium consumption plan

$$c_1^* = (2, 2, 3, 1)$$

is not budget feasible with the sublinear Knight-Walras prices because we have

$$\begin{aligned}\Psi^{\mathcal{P}}(c_1^* - e_1) &= \max_{p\in\left[\frac{1}{8},\frac{3}{8}\right]} \frac{1}{4}(2 - 3 + 2 - 1) + p \cdot (3 - 2) + \left(\frac{1}{2} - p\right) \cdot (1 - 2) \\ &= \max_{p\in\left[\frac{1}{8},\frac{3}{8}\right]} 2p - \frac{1}{2} - p = \frac{1}{4} > 0.\end{aligned}$$

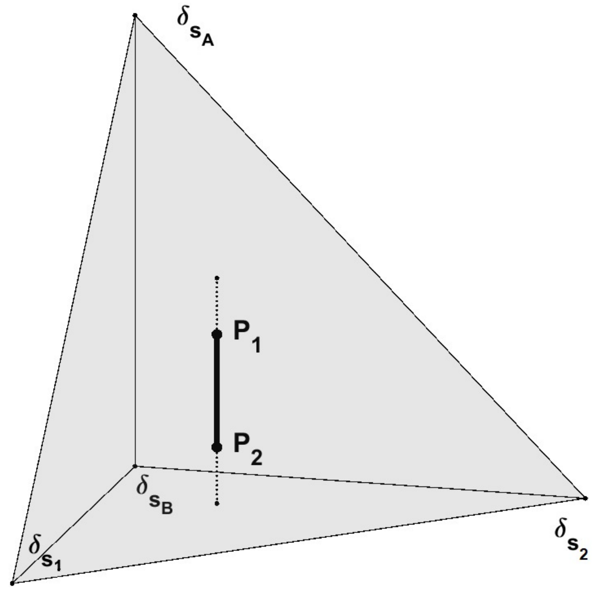

Figure 1: In view of Example 1, price uncertainty based on P_0 (rather than the smaller set of relevant beliefs) is given by $\mathcal{P}_0 = lin(P_1, P_2) \cap \Delta = cx\{(\frac{1}{4}, \frac{1}{4}, \frac{1}{2}, 0), (\frac{1}{4}, \frac{1}{4}, 0, \frac{1}{2})\}$ (dashed line segment), while $\mathcal{P} = cx(P_1, P_2)$ (solid line segment).

The previous theorem demonstrates that a market which appropriately accounts for the heterogeneity of agents' beliefs can prevent unjustified, belief-inefficient trade in equilibrium. In real-world scenarios, however, the market price system—whether it be a market maker or a regulator—may have limited information about the precise subjective beliefs $P_1, \ldots, P_I$ of the agents in the economy. This raises the question of whether it is possible to achieve the same outcome without knowing the exact beliefs of the agents. We will address this question in the following.

Let us write $P_0(A) = P_i(A)$ for the common belief for events $A \in \mathcal{G}$. Let $\mathcal{P}_0 \supset \mathcal{P}$ denote the set of all probability measures on S that extend P_0 from $\mathcal{G}$ to the power set $\mathcal{F}$ of S, see Figure 1 for an illustration within the setting of our leading Example 1. Then the conclusion of Theorem 2 still holds true.

Theorem 3 *Fix any set of priors $\mathcal{Q}$ satisfying $\mathcal{P} \subseteq \mathcal{Q} \subseteq \mathcal{P}_0$. In every Knight-Walras equilibrium (ψ^*, c^*) with respect to $\mathcal{Q}$, the equilibrium allocation c^* is belief-neutral efficient.*

The result shows that any set $\mathcal{Q}$ of priors situated between the set of reasonable beliefs $\mathcal{P}$ and the set $\mathcal{P}_0$ of consistent extensions of a reference measure P_0 yields belief-efficient equilibrium allocations. Next, we explore whether it is possible to use a smaller set $\mathcal{Q}$ than the set of reasonable beliefs $\mathcal{P}$ for the price functional of a Knight-Walras equilibrium while still achieving a belief-neutral efficient allocation. This question is closely linked to the subspace of ambiguity-free or frictionless trades associated with $\mathcal{Q}$. We demonstrate that if the set $\mathcal{Q}$ permits a broader range of ambiguity-free trades than those induced by the field $\mathcal{G}$, belief-neutral efficiency is no longer preserved.

Theorem 4 *Let Q be a probability measure on some sigma algebra $\mathcal{H} \supsetneq \mathcal{G}$, with $Q = P_0$ on $\mathcal{G}$. Define the set of priors by $\mathcal{Q} = \{P \in \Delta S : P = Q \text{ on } \mathcal{H}\}$. In every Knight-Walras equilibrium (ψ^*, c^*) with respect to $\mathcal{Q}$, the equilibrium allocation c^* is belief-neutral inefficient.*

The modified versions of the First Welfare Theorem, as stated in Theorems 2 and 3, hold only when the set of reasonable beliefs a social planner might use and the set of priors a market maker employs to determine prices are sufficiently aligned. Theorem 4 establishes a lower threshold: if price uncertainty $\mathcal{Q}$ is strictly smaller than $\mathcal{P}$, the resulting Knight-Walras equilibrium fails to achieve belief-neutral efficiency. Intuitively, in this scenario, an excess of unjustified bets remains budget-feasible, leading to excessive speculation in equilibrium. Moreover, it is not viable to consider an alternative set that excludes the set of reasonable beliefs $\mathcal{P}$. For instance, incorporating too many priors into the price functional may result in a complete lack of trade in the Knight-Walras equilibrium, which generally does not correspond to a belief-neutral efficient outcome (see Theorem 8 in Beissner and Riedel (2019)).

The set of reasonable beliefs $\mathcal{P} \subset \Delta S$ and the price uncertainty $\mathcal{Q} \subset \Delta S$ represent conceptually distinct elements. The set $\mathcal{P}$ consists of individuals' subjective beliefs. From a normative standpoint, a social planner may be reluctant to designate any single belief as the "correct" one. Instead, they might prefer to rely on the entire set of reasonable beliefs as the foundation for welfare evaluations (see Sections II.C.2 and II.C.3 of Brunnermeier, Simsek, and Xiong (2014)). In contrast, price uncertainty $\mathcal{Q}$ serves as a pricing parameter within the equilibrium framework. It reflects the market maker's degree of caution by imposing constraints on agents' budget sets. For further details and discussion, see Section 2.3 of Beissner and Riedel (2019).

5 Implementation via Transaction Costs

The abstract results presented above do not reveal how a market maker might structure markets to mitigate undesirable speculative trading. In this section, we demonstrate how this objective can be accomplished by designing appropriate financial markets that incorporate transaction costs for purely speculative assets.

Consider the following financial market with $m+1$ assets. The assets are traded at bid–ask prices $0 \le q_j^B \le q_j^A, j = 0, \dots, m$ and payoff $x_j \in \mathbb{X}_+, j = 0, \dots, m$. We furthermore assume that "cash" is frictionless and normalize the interest rate to zero: $q_0^B = q_0^A = x_0 = 1$. We call $\mathcal{M} = (q^A, q^B, x)$ a financial market with transaction costs.

We assume that the consumption good is traded at a price $\phi(s) \ge 0$ on a spot market at time 1 in state s. Agents have to finance their desired value of excess demand $\phi(c_i - e_i)$ by trading in the financial market. A portfolio is a vector $\theta = (\theta^A, \theta^B) \in \mathbb{R}_+^{2(m+1)}$; θ_j^A is the number of assets j bought at time 0, θ_j^B is the number of assets j sold short at time 0. We say that a portfolio θ superhedges a contingent plan $\xi \in \mathbb{X}$ if its cost is less or equal zero,

$$\gamma(\theta) := \sum_{j=0}^{M} \left(\theta_j^A q_j^A - \theta_j^B q_j^B\right) \le 0\,,$$

and its value at time 1 in state s suffices to cover $\xi(s)$, i.e.

$$V(\theta) := \sum_{j=0}^{M} \left(\theta_j^A - \theta_j^B\right) x_j \ge \xi\,.$$

The budget set of agent i in the transaction cost economy $\mathcal{E}^{tc}$ (under $\mathcal{M}$) consists of all consumption plans c_i whose value of excess demand can be financed in the sense that there is $\theta \in \mathbb{R}_+^{2(m+1)}$ with $\gamma(\theta) \le 0$ and $V(\theta) \ge \phi(c_i - e_i)$. Write $B^{tc}(\phi, e^i)$ for the budget set.

A financial equilibrium consists of a spot price $\phi : \Omega \to \mathbb{R}_+$, an allocation $c = (c_i)_{i=1,\dots,I} \in \mathbb{X}_+^I$, and portfolios $\theta_i = \left(\theta_i^A, \theta_i^B\right) \in \mathbb{R}_+^{2(m+1)}, i = 1, \dots, I$ such that[4]

1. the allocation c is feasible, i.e. $\sum_{i=1}^I (c_i - e_i) \le 0$,
2. the financial market clears, i.e. $\sum_{i=1}^I \theta_{ij}^A = \sum_{i=1}^I \theta_{ij}^B$ for $j = 0, \dots, m$,
3. for each agent i, c_i is optimal in the budget set $\mathbb{B}^{tc}(\phi, e_i)$, i.e. $c_i \in \mathbb{B}^{tc}(\phi, e_i)$ and for all $d \in \mathbb{X}_+$ with $U_i(d) > U_i(c_i)$ we have $d \notin \mathbb{B}^{tc}(\phi, e_i)$.

Let us now sketch how one can design a financial market with transaction cost that implements the desired belief-neutral efficient allocation of a Knight-Walras equilibrium[5]. The basic idea goes as follows: We need to find a set of

[4]Note that we require a double index for the portfolios: θ_{ij}^A is the number of assets j that agent i buys at time 0.

[5]The construction is based on our previous previous work Beissner and Riedel (2019), where we have established an equivalence between Knight-Walras equilibria and transaction cost equilibria, and on Araujo, Chateauneuf, and Faro (2018) who provide an excellent analysis of the relation between sublinear expectations and pricing rules generated by markets with transaction costs.

securities with certain payoffs and corresponding bid–ask spreads such that the budget constraint in the Knight–Walras economy and the budget constraint in the financial market are equivalent. One can find the securities that span the financial market by looking at the convex cones of assets x that are priced by the belief of agent i:

$$V_i = \{x \in \mathbb{X} : \mathbb{E}^{\mathcal{P}} x = E^{P_i} x\}.$$

This convex cone is polyhedral because it is the intersection of finitely many half-spaces with normal vectors $P_j - P_i$. By the Minkowski–Weyl duality, we know that it is generated by finitely many payoffs $x_{i,1}, \ldots, x_{i,r_i}$:

$$V_i = \operatorname{cone}(x_{i,1}, \ldots, x_{i,r_i}).$$

We take these payoffs as the dividends for our assets and define the prices

$$q^B_{i,j} = \min_{l=1,\ldots,n} E^P_l x_{i,j}, \quad q^A_{i,j} = \max_{l=1,\ldots,n} E^P_l x_{i,j}, \quad j = 1, \ldots, r_i.$$

We add a riskless security (cash), and are then able to prove the following theorem.

Theorem 5 *Let $(\psi^*, (c_i^*))$ be a Knight-Walras equilibrium with respect to the set of reasonable beliefs $\mathcal{P}$ such that the equilibrium allocation (c_i^*) is belief-neutral efficient. There exists a financial market with transaction costs $\mathcal{M}$ and a financial equilibrium with spot price $\phi = \psi^*$, allocation (c_i^*), and portfolios $\theta_i = \left(\theta_i^A, \theta_i^B\right), i = 1, \ldots, I$.*

The proof is an application of Theorem 9.2 in Beissner and Riedel (2019); we refer to Araujo, Chateauneuf, and Faro (2018) for more details on pricing rules in financial markets with transaction costs. We illustrate the theorem by constructing the suitable financial market with transaction costs in our running example.

Example 1 *(continued) Suppose that next to the riskless asset x_0, we have an asset that pays off the endowment of agent 1 as a dividend, i.e. $x_1 = e_1 = (3, 1, 2, 2)$. Furthermore, there are two Arrow securities for the outcome of the sport event, i.e. $x_2 = (0, 0, 1, 0)$ and $x_3 = (0, 0, 0, 1)$. This financial market is complete as every potential consumption plan can be replicated by trade in the assets.*

If the assets were traded in a frictionless way, the classic equivalence result between Arrow-Debreu and Radner equilibria would lead to belief-neutral inefficient equilibria and the agents would bet on the sports events. However, things change if we introduce transaction costs. Define

$$q^B_j = \min\left(E^{P_1} x_j, E^{P_2} x_j\right), \quad q^A_j = \max\left(E^{P_1} x_j, E^{P_2} x_j\right), \quad j = 1, 2, 3.$$

As the endowment e_1 is $\mathcal{G}$-measurable, the corresponding asset 1 is frictionless, with $q_1^B = q_1^A = 2$. The Arrow securities exhibit transaction costs, with

$$q_j^B = \frac{1}{8} < q_j^A = \frac{3}{8}, \quad j = 2, 3.$$

In this situation, both agents stop speculating on the sports event as they are no longer yielding positive expected profits under either belief. The resulting transaction cost equilibrium leads to the same belief-neutral efficient allocation as the Knight-Walras equilibrium. Further details can be found in the appendix.

The theorem has significant policy implications. Specifically, we explore how the normative selection of an efficiency criterion in the presence of heterogeneous beliefs inherently shapes the role of transaction costs as both market frictions and a regulatory instrument.

The financial market $\mathcal{M}$ from Theorem 5 relies on choosing the right security design. To conclude this section, we list several appealing properties:

1. The market friction does not prevent beneficial trade: there is no bid-ask spread for $\mathcal{G}$-measurable contingent claims. Moreover, all portfolios on the security market line have no bid-ask spreads.

2. The Knight-Walras price system Ψ is a super-replication price of $\mathcal{M}$ and arbitrage free, i.e., for every nonzero $x \in \mathbb{X}_+$ we have $\Psi(x) > 0$.

3. The financial market is efficient complete.[6] The spanning securities $(x_0, x_1, \ldots, x_m)$ are anti-comonotone with regards to the state price ψ^*.

6 Conclusion

Heterogeneous subjective beliefs, regardless of their source, can lead to problematic betting behavior. We investigate how market structures can be designed to facilitate seamless hedging and insurance-based trading while limiting speculative trades driven purely by betting. Our results demonstrate that sublinear pricing, which accurately reflects the ambiguity arising from diverse beliefs, effectively achieves this goal. Moreover, this pricing framework can be implemented through appropriately designed transaction costs.

We replace the classic notion of Pareto efficiency with an alternative belief-neutral efficiency notion. While Pareto efficiency and frictionless Walrasian

[6]To understand property 3., let $L_\Psi = \{x \in \mathbb{X} : y > x \Rightarrow \Psi(y) > \Psi(x)\}$ denote the space of undominated payoffs of some pricing rule $\Psi : \mathbb{X} \to \mathbb{R}$. Based on the notion of efficient securities, see Dybvig (1988), undominated payoffs are efficient securities see Araujo, Chateauneuf, and Faro (2018). Ψ is the pricing rule of an *efficient complete market* if $L_\Psi = \mathbb{X}$.

markets go hand in hand, we show that belief-neutral efficiency requires a different market design, an idea that goes as far back as Hurwicz (1973), maybe. One can thus view our paper as a contribution to *market design*, creating a relation between General Equilibrium Theory and Design, in a similar way as Mechanism Design relates Game Theory and social choice functions. A natural next step in exploring associated financial systems is to ask: What further connections exist between efficiency notions, market mechanisms, and financial market frictions?

A Proof of Theorem 1

We start with the proof of Lemma 1 that we need in the proof of Theorem 1.

Proof of Lemma 1: Recall that a feasible allocation $(c_i^*) \in \Lambda(e)$ is belief-neutral efficient if it is efficient in every homogeneous belief economy $\mathcal{E}(Q, \dots, Q)$ for all reasonable beliefs $Q \in \mathcal{P}$. Fix some $Q \in \mathcal{P}$. By Dana (1993), the allocation (c_i^*) maximizes the linear social welfare function

$$\sum_{i=1}^{I} \alpha_i E^Q u_i(d_i)$$

over all feasible allocations $(d_i) \in \Lambda(e)$ for some weights $\alpha \in \Delta I$. By linearity of the (homogeneous) expectation, one can maximize the sum

$$\sum_{i=1}^{I} \alpha_i u_i(d_i(s))$$

pointwise for every state s, yielding the first-order conditions

$$\alpha_i u_i'(c_i^*(s)) = \alpha_j u_j'(c_j^*(s)) \tag{5}$$

for all agents $i, j \in \{1, \dots, I\}$ and all states $s \in S$. These equations are independent of the chosen $Q \in \mathcal{P}$, and we thus conclude that (c_i^*) is efficient in $\mathcal{E}(P, \dots, P)$ for every $P \in \mathcal{P}$.

The other direction follows similarly. If we have the first-order condition (5), then the necessary and sufficient conditions for a feasible allocation to maximize a weigthed social welfare function are satisfied, and (c_i^*) is efficient in every homogeneous belief economy $\mathcal{E}(P, \dots, P)$ for $P \in \mathcal{P}$. $\square$

Let us now turn to a rigorous formulation of Theorem 1 and its proof. We parametrize the set of economies by the belief profiles of the form $\mathfrak{B} = (P_1, \dots, P_I)$, keeping the other ingredients of the economy fixed.

A belief P with full support on $S = \{s_1, \ldots, s_n\}$ can be equivalently described by a strictly positive vector $(P(s_1), \ldots, P(s_{n-1})) \in \mathbb{R}^{n-1}$ with $\sum_{k=1}^{n-1} P(s_k) < 1$. Denote this open set by $\mathbb{B} \subset \mathbb{R}^{n-1}$. Our set of economies

$$\mathfrak{E} = \{\mathcal{E}(P_1, \ldots, P_I) : \text{ all } P_i \text{ have full support}\}$$

is thus isomorphic to $\mathbb{B}^I \subset \mathbb{R}^{(n-1)I}$.

Theorem A.1 *The set of economies $\mathcal{E}(P_1, \ldots, P_I)$ that admit belief-neutral efficient Arrow-Debreu equilibria is both a Lebesgue null set and nowhere dense in $\mathfrak{E}$.*

PROOF: Fix a belief profile $\mathfrak{B} = (P_1, \ldots, P_I)$. Let $\left(\Psi, \left(c_i^{\mathfrak{B}}\right)\right)$ be an Arrow-Debreu equilibrium in the economy $\mathcal{E}(\mathfrak{B})$. By the first welfare theorem, the allocation $\left(c_i^{\mathfrak{B}}\right)$ is efficient. It is an interior allocation because initial endowments are strictly positive and utilities strictly monotone and satisfy the Inada condition. Note that the utility $U_i(c) = E^{P_i} u_i(c)$ can be written as an expected utility under a common probability P with full support and state-dependent expected utility:

$$U_i(c) = \sum_{s \in S} v_i(s, c(s)) P(s)$$

with

$$v_i(s, c) = \frac{P_i(s)}{P(s)} u_i(c).$$

By Theorem 1 in Dana (1993), there exist strictly positive weights $\alpha_i > 0$ such that we have

$$\alpha_i \frac{P_i(s)}{P(s)} u_i'(c_i(s)) = \alpha_j \frac{P_j(s)}{P(s)} u_j'(c_j(s)). \tag{6}$$

Now assume that the allocation $\left(c_i^{\mathfrak{B}}\right)$ is belief-neutral efficient. From Lemma 1, we conclude that there exist weights $\lambda_i > 0$ such that we have the belief-independent characterization

$$\lambda_i u_i'(c_i(s)) = \lambda_j u_j'(c_j(s)). \tag{7}$$

From equations (6) and (7), $P_1(s) = P_2(s)$ follows. As i and j have been arbitrary, we conclude that efficient and belief-neutral efficient Arrow-Debreu equilibria exist only in homogeneous belief economies.

The next lemma then concludes. □

Lemma A.2 *The set of homogeneous belief economies is a Lebesgue null set and nowhere dense in the set of economies with possibly heterogeneous beliefs.*

Proof: The set of homogeneous beliefs economies is given by the requirement $P_1 = P_2 = \ldots = P_I$, thus the intersection of a $(n-1)$-dimensional plane in the space $\mathbb{R}^{(n-1)I}$ with the open set $\mathbb{B}^I$, so a Lebesgue null set. The requirement $P_1 = P_2 = \ldots = P_I$ also yields a closed set in the relative topology of $\mathbb{B}^I$ whose interior is empty, hence the set is also nowhere dense. $\square$

B Proof of Theorem 2 and 3

Proof of Theorem 2: Fix any $P \in \mathcal{P}$, and consider the homogeneous belief economy $\mathcal{E}(P, \ldots, P)$. Recall the definition of the representative agent's Bernoulli utility function u_α in (2) for homogeneous beliefs economies and the resulting allocation rule $(C_{\alpha,i})$ in (3). By Theorem 3.1 of Dana (1993), the equilibrium is characterized as follows. There is a vector of Negishi weights $\alpha \in \Delta I$ such that $(\Psi, (c_i^*))$ is an Arrow-Debreu equilibrium in $\mathcal{E}(P, \ldots, P)$ with $\Psi(x) = E^P[\psi x]$ and

$$\psi(s) = u'_\alpha(e(s)) \tag{8}$$

and

$$c_i^*(s) = C_{\alpha,i}(e(s)).$$

As all initial endowments are strictly positive and we have Inada's conditions for the Bernoulli utilities, the allocation is an interior allocation and all weights α_i are strictly positive. By Lemma 1, the allocation (c_i^*) is belief-neutral efficient. The state price ψ and the equilibrium consumption plans c_i^* are functions of aggregate endowment e, and thus $\mathcal{G}$–measurable.

We now show that $(\psi, (c_i^*))$ is a Knight-Walras equilibrium with respect to $\mathcal{P}$, using Theorem 4 in Beissner and Riedel (2019). We first need to show that c_i^* is budget-feasible for agent i according to the Knight-Walras budget constraint. This follows immediately from the fact that c_i^* is budget-feasible in the Arrow-Debreu equilibrium of the homogeneous beliefs economy:

$$\Psi^{\mathcal{P}}(c_i^* - e_i) = \max_{Q \in \mathcal{P}} E^Q[\psi(c_i^* - e_i)] = E^P[\psi(c_i^* - e_i)] \leq 0. \tag{9}$$

where the second equality follows from the $\mathcal{G}$-measurability of the net trade. We next have to check that the set of subjective beliefs $\pi_i(c_i^*)$ and the set of effective pricing measures

$$\phi^{\mathcal{P}}(c_i^* - e_i) = \left\{ Q \in \Delta S : E^Q[c_i^* - e_i] \geq E^Q \eta \text{ for all } \eta \text{ with } \Psi^{\mathcal{P}}(\eta) \leq \Psi^{\mathcal{P}}(c_i^* - e_i) \right\}$$

have a nonempty intersection. So let us compute those sets in our special case.

For subjective expected utility, the subjective belief $\pi_i(x)$ at $x \in \mathbb{X}$ is a singleton and it consists of the so-called risk-adjusted prior

$$\pi_i(x)(s) = \frac{P_i(s) u'_i(x(s))}{\sum_{t \in S} P_i(t) u'_i(x(t))}, \tag{10}$$

compare Section 2.4 in Rigotti, Shannon, and Strzalecki (2008) (subjective expected utility is a special case of maxmin preferences).

According to Proposition 2 in Beissner and Riedel (2019), elements in $\phi^{\mathcal{P}}(c_i^* - e_i)$ take the form

$$\frac{\psi(s)Q(s)}{\sum_{t\in S}\psi(t)Q(t)},$$

for some $Q \in \arg\max_{Q\in\mathcal{P}} E^Q[\psi(c_i^* - e_i)]$. Now, the state sprice ψ, the consumption plan c_i^*, and the endowment e_i are $\mathcal{G}$-measurable; as the beliefs in $\mathcal{P}$ all coincide on $\mathcal{G}$, we have $E^Q[\psi(c_i^* - e_i)] = E^{Q'}[\psi(c_i^* - e_i)]$ for all $Q, Q' \in \mathcal{P}$. Hence, we conclude that the set of effective pricing measures consists of measures of the form

$$\frac{\psi(s)Q(s)}{\sum_{t\in S}\psi(t)Q(t)},$$

for some $Q \in \mathcal{P}$. From (3) and (8, we conclude that we can take $Q = P_i$ to see that $\pi_i(c_i^*) \in \phi^{\mathcal{P}}(c_i^* - e_i)$. Hence, $(\psi, (c_i^*))$ is a Knight-Walras equilibrium with respect to $\mathcal{P}$. □

Proof of Theorem 3: Inspecting the proof of Theorem 2, we see that we can replace $\mathcal{P}$ by $\tilde{\mathcal{P}}$. This larger price uncertainty is consistent with the set of reasonable beliefs on $\mathcal{G}$, since still contained in $\mathcal{P}_0$. This again leads to $\mathcal{G}$ measurability of the belief neutral efficient candidate allocation (c_i^*) and state price ψ^*. More precisely, we get again (8), i.e., $(\psi^*, (c_i^*))$ is again a continuous function of e. All further arguments apply too, since the following arguments solely depend on the $\mathcal{G}$-measurability of ψ^*, c_i^*, e_i and $\psi^*(c_i^* - e_i)$. □

C Proof of Theorem 4

To begin, recall the risk-adjusted prior $\pi_i(x)$ of agent i at x defined in (10). For the Knight-Walras equilibrium allocation in Theorem 2, $\pi_i(c_i^*) \neq \pi_j(c_j^*)$ holds for some i, j. We prepare the proof with the following lemma.

Lemma C.1 *Let (c_i^*) be a Knight-Walras equilibrium allocation with respect to $\mathcal{Q}$, where $\mathcal{Q}$ be as in Theorem 4. There is a belief π on $\mathcal{H}$ such that*

$$\pi \;=\; E\,[\,\pi_i(c_i^*)\,|\mathcal{H}] \qquad \textit{for all } i, \tag{11}$$

where E denotes the expectation under the uniform distribution on S.

Proof of Theorem 4: By assumption, $\mathcal{G}$ is the maximal domain where the agent's belief agree. Hence, the sub σ-algebra $\mathcal{H}$ is strictly finer than $\mathcal{G}$, such that the agents' beliefs are non-concordant with respect to $\mathcal{H}$, i.e., $E[P_k|\mathcal{H}]$ is not $\mathcal{G}$-measurable for some k and $E[P_k|\mathcal{H}] \neq E[P_i|\mathcal{H}]$ for some i, k. This also implies that $\mathbb{L}^{\mathcal{P}_0} \subsetneq \mathbb{L}^{\mathcal{Q}}$ holds, by the definition of $\mathcal{P}_0$ and $\mathcal{Q}$.

On the other hand, the price uncertainty $\mathcal{Q}$ consists of all consistent extensions of $Q : \mathcal{H} \to [0,1]$ to all of $\mathcal{F}$. Let

$$\mathbb{L}(\mathcal{H}) = \{x \in \mathbb{X} : x \text{ is } \mathcal{H}\text{-measurable}\}.$$

The present form of price uncertainty implies that $\mathbb{L}^{\mathcal{Q}} = \mathbb{L}(\mathcal{H})$. We show in the following claim that there is trade based on belief disagreement on the "events" in $\mathcal{H} \setminus \mathcal{G}$, where the reduced price uncertainty $\mathcal{Q}$, relative to $\mathcal{P}_0$, fails to mute trade.

Claim: There is an agent k, such that c_k^* is not $\mathcal{G}$-measurable.
Proof: Set the constant $\mu_i = \sum_{s \in S} P_i(s) u_i'(c_i^*(s)) > 0$. Recall the form of $\pi_i(c_i^*)$ defined in (10) We apply Lemma C.1, the $\mathcal{H}$-measurability of each c_i^* and derive:

$$\begin{aligned} E[\pi_i(c_i^*)|\mathcal{H}] = E[\pi_k(c_k^*)|\mathcal{H}] \quad &\Leftrightarrow \quad E\left[P_i \frac{u_i'(c_i^*)}{\mu_i} \Big| \mathcal{H}\right] = E\left[P_k \frac{u_k'(c_k^*)}{\mu_k} \Big| \mathcal{H}\right] \\ &\Leftrightarrow \quad \frac{u_i'(c_i^*)}{\mu_i} E[P_i|\mathcal{H}] = \frac{u_k'(c_k^*)}{\mu_k} E[P_k|\mathcal{H}] \\ &\Leftrightarrow \quad \frac{u_i'(c_i^*)\mu_k}{u_k'(c_k^*)\mu_i} = \frac{E[P_k|\mathcal{H}]}{E[P_i|\mathcal{H}]} \end{aligned}$$

As $\frac{E[P_k|\mathcal{H}]}{E[P_i|\mathcal{H}]}$ fails to be $\mathcal{G}$-measurable by the assumption that $\mathcal{G}$ is the largest domain of common agreement of probabilities, it follows, that either c_k^* or c_i^* fails to be $\mathcal{G}$-measurable, since $x \mapsto u_k'(x)$ and $x \mapsto u_i'(x)$ are strictly decreasing functions, and $\frac{\mu_k}{\mu_i}$ is a constant. This proves the claim.

To complete the proof, assume that c^* is belief–neutral efficient. Note that each c_i^* is $\mathcal{H}$-measurable. By the Claim, there is some agent k where $c_k^* \notin \mathbb{L}(\mathcal{G})$. It then follows by Lemma 1 in view of (3) that the equilibrium allocation satisfies the following functional dependency:

$$c_i^* = C_{\alpha,i}(e) = (u_i')^{-1}\left(\alpha_i u_\alpha'(e)\right) \quad \text{for each } i.$$

But since e is $\mathcal{G}$-measurable by assumption and $x \mapsto C_{\alpha,i}(x)$ is continuous, it follows that each c_i^* is also $\mathcal{G}$-measurable, a contradiction. $\square$

Proof of Lemma C.1: Set $L := \mathbb{L}^{\mathcal{R}}$, where $\mathcal{R} = \{\frac{\psi}{|\psi|} P : P \in \mathcal{Q}\} \subset \Delta S$.

We first show that

$$c^* \in \Lambda_0(e) := \{(x_i) \in \Lambda(e) : (x_i - e_i) \in L\}.$$

By Theorem 6 of Beissner and Riedel (2019), c^* satisfies $\psi(c_i^* - e_i) \in \mathbb{L}^{\mathcal{Q}}$ for all i. This is equivalent to $c_i^* - e_i \in L$, as Lemma C.2 implies $L = \mathbb{L}(\mathcal{H}) = \mathbb{L}^{\mathcal{Q}}$. Since c^* is a Knight-Walras equilibrium allocation, we have $c^* \in \Lambda(e)$ is uncertainty neutral efficient thus constrained efficient with respect to L. We apply Corollary 12.6 in Magill and Quinzii (2002) and get

$$\mathrm{pr}_L(\nabla U_i(c_i^*)) = \mathrm{pr}_L(\nabla U_j(c_j^*)) \text{ for all } i, j, \tag{12}$$

where $\mathrm{pr}_L : \mathbb{X} \to L$ denptes the orthogonal projection onto L with respect to $\langle x, y\rangle = E[xy]$.

We show that L is a sublattice of $\mathbb{X}$. As $(\psi^*, (c_i^*))$ is a Knight-Walras equilibrium with respect to $\mathcal{Q}$, ψ^* is $\mathcal{H}$-measurable and it follows by Lemma C.2 that $L = \mathbb{L}(\mathcal{H})$, since $\mathbb{L}^{\mathcal{Q}} = \mathbb{L}(\mathcal{H})$. Since L is a lattice,[7] the projection pr_L coincides with the conditional expectation $E[\cdot|\mathcal{H}] : \mathbb{X} \to \mathbb{L}(\mathcal{H})$. Thus, (12) takes the form in (11). □

The proof of Lemma C.1 employs the following result.

Lemma C.2 *For any $\mathcal{G}$-measurable $\psi : \Omega \to \mathbb{R}_{++}$, we have*

$$\mathbb{L}^{\mathcal{Q}_0} = \mathbb{L}^{\mathcal{P}_0}, \quad \textit{where } \mathcal{Q}_0 := \left\{ \frac{\psi}{|\psi|} P : P \in \mathcal{P}_0 \right\} \subset \Delta S.$$

Proof of Lemma C.2: Consider the capacity defined by

$$\mu_{P_0}(A) = \max_{B \subseteq A, B \in \mathcal{G}} P_0(B), \tag{13}$$

where $P_0 \in \Delta\mathcal{G}$ is the belief on $\mathcal{G}$, where the agents agree. It can be shown that μ_{P_0} is a convex capacity with core

$$\mathrm{core}(\mu_{P_0}) = \{Q \in \Delta S : Q(A) \geq \mu_{P_0}(A) \text{ for all } A \in \mathcal{F}\}.$$

A measure Q is in the core if and only if it extends P_0 to $\mathcal{F}$, i.e., $\mathrm{core}(\mu_{P_0}) = \mathcal{P}_0$.

Define $Q_0 = \frac{\psi}{|\psi|} P_0$, since ψ is $\mathcal{G}$-measurable, we have $Q_0 \in \Delta\mathcal{G}$. As in (13) for $P_0 \in \Delta\mathcal{G}$, Q_0 gives rise to define a new inner capacity $\mu_{Q_0}(A) = \max_{B \subset A, B \in \mathcal{G}} Q_0(B)$. This in turn allows for a characterization of $\mathcal{Q}_0$ via

$$\mathcal{Q}_0 = \mathrm{core}(\mu_{Q_0}) = \{Q \in \Delta S : Q(G) = Q_0(G),\ \forall G \in \mathcal{G}\}.$$

The claim follows since $\mathbb{L}^{\mathcal{Q}_0} = \mathbb{L}(\mathcal{G})$ holds by the same reasoning that proves the equality $\mathbb{L}^{\mathcal{P}_0} = \mathbb{L}(\mathcal{G})$. □

[7] Given an orthogonal projection $p : \mathbb{L}(\mathcal{F}) \to L$ for some (closed) sub-vector space L, the following is equivalent: (i) p is a conditional expectation. (ii) $L = \mathbb{L}(\mathcal{H})$ for some σ-algebra $\mathcal{H}$. See Theorem 22.5 in Schilling (2017) for the general case.

D Details for the Knight-Walras equilibrium of the sports example

We need to show that the allocation $c_i^*(s) = 2$ with spot price $\psi(s) = 1$ in all states is a Knight–Walras equilibrium for the price functional

$$\Psi^{\mathcal{P}}(x_1, x_2, x_A, x_B) = \frac{x_1 + x_2 + x_A + x_B}{4} + \frac{1}{8}(x_A - x_B)^+ + \frac{1}{8}(x_B - x_A)^+ .$$

The market obviously clears and the budget constraint is satisfied because

$$\Psi^{\mathcal{P}}(c_1^* - e_1) = \Psi^{\mathcal{P}}(1, -1, 0, 0) = 0$$

and

$$\Psi^{\mathcal{P}}(c_2^* - e_2) = \Psi^{\mathcal{P}}(-1, 1, 0, 0) = 0.$$

It remains to show optimality for c_i^* given the Knight–Walras budget constraint. So suppose that $\Psi^{\mathcal{P}}(c_i - e_i) \leq 0$. Then we have by concavity and the fact that $c_i^* = 2$ is constant and $P_i \in \mathcal{P}$

$$\begin{aligned}
\mathbb{E}^{P_i} u_i(c_i) - \mathbb{E}^{P_i} u_i(c_i^*) &\leq u_i'(2)\mathbb{E}^{P_i}(c_i - c_i^*) \\
&= u_i'(2)\left(\mathbb{E}^{P_i}(c_i - e_i) + \mathbb{E}^{P_i}(e_i - 2)\right) \\
&\leq u_i'(2)\left(\Psi^{\mathcal{P}}(c_i - e_i) + 0\right) \\
&\leq u_i'(2)\,(0 + 0) = 0.
\end{aligned}$$

We conclude that c_i^* is indeed optimal.

One can even show that this equilibrium is the unique Knight–Walras equilibrium for this economy. We leave the proof to the reader.

References

Araujo, A., A. Chateauneuf, and J. H. Faro (2018): "Financial market structures revealed by pricing rules: Efficient complete markets are prevalent," *Journal of Economic Theory*, 173, 257–288.

Athanasoulis, S. G., and R. J. Shiller (2000): "The significance of the market portfolio," *The Review of Financial Studies*, 13(2), 301–329.

Banerjee, S., and I. Kremer (2010): "Disagreement and learning: Dynamic patterns of trade," *The Journal of Finance*, 65(4), 1269–1302.

Beissner, P., and F. Riedel (2019): "Equilibria Under Knightian Price Uncertainty," *Econometrica*, 87(1), 37–64.

Bénabou, R., and J. Tirole (2016): "Mindful economics: The production, consumption, and value of beliefs," *Journal of Economic Perspectives*, 30(3), 141–164.

Blume, L. E., T. Cogley, D. A. Easley, T. J. Sargent, and V. Tsyrennikov (2018): "A case for incomplete markets," *Journal of Economic Theory*, 178, 191–221.

Borovička, J. (2020): "Survival and long-run dynamics with heterogeneous beliefs under recursive preferences," *Journal of Political Economy*, 128(1), 206–251.

Brunnermeier, M., A. Simsek, and W. Xiong (2014): "A Welfare Criterion For Models With Distorted Beliefs," *The Quarterly Journal of Economics*, 129(4), 1753–1797.

Brunnermeier, M. K., and J. A. Parker (2005): "Optimal expectations," *American Economic Review*, 95(4), 1092–1118.

Dana, R. (1993): "Existence and uniqueness of equilibria when preferences are additively separable," *Econometrica*, 61(4), 953–957.

Danan, E., T. Gajdos, B. Hill, and J.-M. Tallon (2016): "Robust social decisions," *American Economic Review*, 106(9), 2407–25.

Daniel, K. D., D. Hirshleifer, and A. Subrahmanyam (2001): "Overconfidence, arbitrage, and equilibrium asset pricing," *The Journal of Finance*, 56(3), 921–965.

Dávila, E. (2023): "Optimal financial transaction taxes," *The Journal of Finance*, 78(1), 5–61.

Detemple, J., and S. Murthy (1994): "Intertemporal asset pricing with heterogeneous beliefs," *Journal of Economic Theory*, 62(2), 294–320.

Dreze, J. H. (1972): "Market allocation under uncertainty," *European Economic Review*, 2(2), 133–165.

Dybvig, P. H. (1988): "Distributional Analysis of Portfolio Choice," *The Journal of Business*, 61(3), 369–93.

Easley, D., and M. O'Hara (2009): "Ambiguity and nonparticipation: The role of regulation," *The Review of Financial Studies*, 22(5), 1817–1843.

Ghirardato, P., F. Maccheroni, and M. Marinacci (2004): "Differentiating ambiguity and ambiguity attitude," *Journal of Economic Theory*, 118(2), 133–173.

GILBOA, I., L. SAMUELSON, AND D. SCHMEIDLER (2014): "No-Betting-Pareto Dominance," *Econometrica*, 82(4), 1405–1442.

GUL, F., AND W. PESENDORFER (2014): "Expected uncertain utility theory," *Econometrica*, 82(1), 1–39.

HAMMOND, P. J. (1981): "Ex-ante and ex-post welfare optimality under uncertainty," *Economica*, 48(191), 235–250.

HARRIS, M., AND A. RAVIV (1993): "Differences of opinion make a horse race," *The Review of Financial Studies*, 6(3), 473–506.

HARRIS, R. (1978): "Ex-post efficiency and resource allocation under uncertainty," *The Review of Economic Studies*, 45(3), 427–436.

HARRISON, J. M., AND D. M. KREPS (1978): "Speculative investor behavior in a stock market with heterogeneous expectations," *The Quarterly Journal of Economics*, 92(2), 323–336.

HURWICZ, L. (1973): "The design of mechanisms for resource allocation," *The American Economic Review*, 63(2), 1–30.

JOUINI, E., AND C. NAPP (2007): "Consensus consumer and intertemporal asset pricing with heterogeneous beliefs," *The Review of Economic Studies*, 74(4), 1149–1174.

JUDD, K. L., F. KUBLER, AND K. SCHMEDDERS (2003): "Asset trading volume with dynamically complete markets and heterogeneous agents," *The Journal of Finance*, 58(5), 2203–2217.

KANDEL, E., AND N. D. PEARSON (1995): "Differential interpretation of public signals and trade in speculative markets," *Journal of Political Economy*, 103(4), 831–872.

MAGILL, M., AND M. QUINZII (2002): *Theory of incomplete markets*, vol. 1. MIT press.

MORRIS, S. (1995): "The common prior assumption in economic theory," *Economics & Philosophy*, 11(2), 227–253.

ODEAN, T. (1998): "Volume, volatility, price, and profit when all traders are above average," *The Journal of Finance*, 53(6), 1887–1934.

——— (1999): "Do investors trade too much?," *American economic review*, 89(5), 1279–1298.

Posner, E., and E. G. Weyl (2013): "Benefit-cost analysis for financial regulation," *American Economic Review*, 103(3), 393–97.

Rigotti, L., C. Shannon, and T. Strzalecki (2008): "Subjective beliefs and ex ante trade," *Econometrica*, 76(5), 1167–1190.

Scheinkman, J. A., and W. Xiong (2003): "Overconfidence and speculative bubbles," *Journal of political Economy*, 111(6), 1183–1220.

Schilling, R. L. (2017): *Measures, integrals and martingales.* Cambridge University Press.

Simsek, A. (2013): "Speculation and risk sharing with new financial assets," *The Quarterly Journal of Economics*, 128(3), 1365–1396.

Starr, R. M. (1973): "Optimal production and allocation under uncertainty," *The Quarterly Journal of Economics*, 87(1), 81–95.

Stiglitz, J. E. (1989): "Using tax policy to curb speculative short-term trading," *Regulatory Reform of Stock and Futures Markets: A Special Issue of the Journal of Financial Services Research*, pp. 3–17.

von Weizsäcker, C. (1969): "General Equilibrium and Uncertainty," *(MIT teaching note).*