Social Choice and Welfare
https://doi.org/10.1007/s00355-023-01467-1

ORIGINAL PAPER

On efficiency in disagreement economies

Christian Heyerdahl-Larsen[1] · **Johan Walden**[2]

Received: 1 June 2022 / Accepted: 8 May 2023

Abstract

We analyze multiple-beliefs based efficiency measures in economies with risk and disagreement, including belief neutral efficiency and inefficiency, incomplete knowledge efficiency, efficiency based on unanimity, and utility aggregators that minimize Bergson welfare functions over multiple beliefs. We provide equivalence results under technical conditions that are satisfied in several work-horse economies, including the exchange economy and a standard economy with a linear production technology. We also provide several examples for which these measures differ. Our results show that the further away one gets from the standard exchange economy, the more the different multiple-beliefs based measures differ in the allocations they identify as efficient, in general. Consequently, the more important the choice of efficiency measure becomes.

1 Introduction

Several recent studies analyze the challenges of using traditional welfare measures in markets with risk and disagreement, see Brunnermeier et al. (2014), Gayer et al. (2014), Blume et al. (2018), and Heyerdahl-Larsen and Walden (2022). Briefly, when agents disagree, speculative trading motives may end up having a prominent role in determining the market outcome, especially in complete markets with severe disagreement. Such outcomes are—via the welfare theorems—efficient in the traditional sense, in that they maximize agents' ex ante expected utilities. But, as these recent studies argue, there is something spurious with such an efficiency measure based on ex ante expected utilities, for which each agent's expectation is based on his/her individual belief. Under disagreement, two risk averse agents who before

We thank Bernard Dumas, Alex Edmans, Willie Fuchs, Stefano Giglio, Brett Green, Lars Hansen, and Christine Parlour, for helpful comments and suggestions. This research was carried out while Christian Heyerdahl-Larsen was affiliated with Indiana University, Kelley School of Business.

✉ Christian Heyerdahl-Larsen
christian.heyerdahl-larsen@bi.no

1 BI Norwegian Business School, Oslo, Norway

2 Haas School of Business, University of California at Berkeley, Berkeley, CA, USA

Published online: 07 July 2023

trading face no risk may end up speculating heavily against each other, both agreeing that one of them ex post will end up very poor, but both believing it will be the other agent. From the social planner's perspective, these speculative trades takes the economy from one with no ex post allocation risk to one with significant risk, i.e., it moves in the opposite direction of what the risk sharing trades in financial markets achieve in an economy with agreement. Indeed, a major purpose of financial markets is to allow for such risk sharing.

The aforementioned studies propose alternatives to the traditional efficiency measure, that are similar in spirit but different in detail. Brunnermeier et al. (2014) define the concepts of belief neutral efficiency and inefficiency. These measures require the social planner to consider efficiency of an allocation for a whole set of beliefs, which includes individual agents' beliefs. For each of these beliefs, efficiency in the traditional sense—using a common belief across agents—is studied. Heyerdahl-Larsen and Walden (2022) introduce a similar measure, in this paper denoted incomplete knowledge efficiency, under the assumption that the social planner does not know which belief within the set is correct. Gayer et al. (2014) define unanimity Pareto dominance, which requires a reallocation to be an improvement under each agent's belief. Blume et al. (2018) use a utility aggregator, e.g., minimizing a Bergson welfare function, over a whole set of admissible beliefs. Common for these novel multiple-beliefs based measures is that they, counter-factually, use common beliefs across agents when comparing allocations, but then take all agents' individual beliefs into account, i.e., use a whole set of such common beliefs, when evaluating efficiency and dominance. We call such efficiency measures multiple-beliefs based.

In this paper, we provide a systematic analysis and comparison of multiple-beliefs based efficiency measures. As mentioned, these measures differ in the technical details of how they are defined. A general takeaway of our analysis is that the measures agree on the set of efficient allocations in some important special cases, but in general they differ. Belief neutral efficiency, belief neutral inefficiency, and incomplete knowledge efficiency, and the maximal solutions to the planner's problem in Blume et al. (2018) are equivalent in an exchange economy environment.[1] In an economy without disagreement, these measures also agree with the so-called U-efficiency, a concept closely related to the unanimity Pareto dominance concept introduced in Gayer et al. (2014). Incidentally, many of the examples provided in the aforementioned literature are for exchange economies, which could wrongly give the impression that the technical differences are not important.

The further away one moves from the exchange economy, the less the measures agree. The unanimity efficient set in general contains speculative outcomes, which the other measures do not. The set of belief neutral efficient allocations may be empty. Indeed, an empty belief neutral efficient set may be viewed as the generic

[1] In general, there may be a "gap" between belief neutral efficiency and inefficiency, i.e., there may be allocations that are neither belief neutrally inefficient, nor belief neutrally efficient. So these two concept are in general not equivalent.

outcome in a production economy setting. The set of maximal solutions in Blume et al. (2018) can in general not be nested with respect to the other concepts.

Incomplete knowledge efficiency, and belief neutral efficiency and inefficiency, are equivalent however, under technical conditions. Specifically, convexity of the utility possibility set and the so-called *reasonable belief set*, together with a strict dominance condition on Pareto efficient allocations under agreement, ensure that the concepts are equivalent. These technical conditions are satisfied in several workhorse production economies. We also provide several examples of economies in which these concepts differ, providing further intuition about why the technical differences matter.

Finally, we provide an example that takes us even further away from the exchange economy setting, in which multiple-beliefs based measures arguably are not meaningful. Altogether, our analysis highlights the challenges of using a multiple-beliefs based approach to define efficiency in economies with disagreement, but also show that these challenges are overcome in important special cases. Our results also raise interesting questions for future research.

Our general aim is not to evaluate or rank the different measures with respect to their usefulness, but rather to point out their similarities and differences, and to understand what drives these differences, and the implications in different economic environments. When we encounter environments in which the applicability of a specific measure is severely limited, we do point this out, however.

Our paper also relates to the literature on beliefs, efficiency, and portfolio choice. Dybvig (1988) studies portfolio efficiency with respect to expenditure minimization. Inefficient portfolios in his model are those that can be replicated in distribution by a less costly portfolio.[2] In our framework, agents will always choose efficient portfolios given their beliefs. Following Mongin (1995), see also Mongin (2016), a large body of literature has analyzed the general aggregation of agent preferences under heterogeneous beliefs, see, e.g., Gilboa et al. (2004) and Fleurbaey (2010).[3] This literature is outside the scope of our paper, which focuses on comparing the recently introduced multiple-beliefs based efficiency measures, as discussed above.

In the next section we introduce the model, and the different efficiency measures. In Sect. 3, we analyze and compare the different measures, and in Sect. 4 we discuss several examples, and draw some general conclusions from our analysis. All proofs are delegated to an Appendix.

[2] Jouini and Kallal (2001) extend the framework of Dybvig (1988) to settings with frictions using risk neutral probabilities.

[3] Bahel and Sprumont (2020) study the class of strategyproof social choice functions when agents have subjective expected utility. Similar to the literature mentioned above, we are not concerned with the incentive-compatibility issue.

2 Model

In this section we present the economic environment and formally define the various measures of efficiency.

Consider an economy with $T \geq 1$ dates, $\mathcal{T} = \{1, 2, \ldots, T\}$, $M \geq 2$ states, and $N \geq 2$ agents.[4] Agents agree on the state space $\Omega = \{\omega_m\}_{m=1,\ldots,M}$ and the filtration $\underline{\mathcal{F}} = (\mathcal{F}_0, \mathcal{F}_1, \ldots, \mathcal{F}_T)$, $\mathcal{F}_0 = \{\emptyset, \Omega\}$, $\mathcal{F} \subset \mathcal{F}_{t+1}$ for all t, and $\mathcal{F}_T = 2^{\Omega}$, but disagree on the probability for different events. Specifically, agent n's beliefs are captured by the probability measure, $\mathcal{Q}^n : \mathcal{F}_T \to [0, 1]$. We define $q^n_m = \mathcal{Q}^n(\{\omega_m\})$ and the M-vectors $q^n = (q^n_1, q^n_2, \ldots, q^n_M) \in S^M$, $n = 1, \ldots, N$. Agents' probability measures are equivalent, $q^n_m > 0 \Leftrightarrow q^{n'}_m > 0$, for all n, n'. For simplicity, we disregard states all agents agree are impossible, so that $q^n_m > 0$ for all n and m, i.e., $q^n \in S^M$ for all n. Here, S^M is the M-dimensional open unit simplex, $S^M \stackrel{\text{def}}{=} \{x \in \mathbb{R}^M_{++} : \sum_m x_m = 1\}$.[5] Agents' beliefs are then summarized in the tuple $\mathbf{q} = (q^1, q^2, \ldots, q^N) \in \left(S^M\right)^N$. We denote the closure of S^M by $\bar{S}^M$.

In the special case when $q^n = q \in S^M$ for all n, this is an *agreement economy*. Whenever $q^n_m \neq q^{n'}_m$ for some n, n', and m, we are in a *disagreement economy*. Agents are said to agree on events in $\mathcal{F}_X \subset \mathcal{F}_T$, if $\mathcal{Q}^n(X) = \mathcal{Q}^{n'}(X)$ for all n, n', and events $X \in \mathcal{F}_X$. Moreover, agents are said to agree that $\mathcal{F}_X$ and $\mathcal{F}_Y$ are independent, if $\mathcal{F}_X$ and $\mathcal{F}_Y$ are independent under each agent's beliefs.

There is one consumption good that agents derive utility from consuming, and a set of production technologies to choose from, with associated allocations of the good between agents. Specifically, the set of nonnegative adapted processes is $\mathcal{L} = \{\ell : \mathcal{T} \times \Omega \to \mathbb{R}^N_+\}$.[6] A nonempty compact set, $\mathcal{A} \subset \mathcal{L}$, denoted the *set of feasible allocations*, determines the joint production and allocation of the consumption good among the agents. Here, for $a \in \mathcal{A}$, $a_{t,m,n} = (a(t, m))_n$ represents the allocation to agent n of the good at time t in state m. We may view $\mathcal{A}$ as a subset of $\mathbb{R}^{T \times M \times N}$. We denote this economic environment, as defined above, the *general economy*.

We will mostly focus on economies that allow for transfers. We define the mapping $\mathcal{P} : \mathbb{R}^{T \times M \times N}_+ \to \mathbb{R}^{T \times M}_+$, such that $X_{m,t} = \mathcal{P}(a) = \sum_n a_{t,m,n}$, represents aggregate production for allocation a at time t in state m, and also the set $\mathcal{A}_X = \mathcal{P}(\mathcal{A}) \subset \mathbb{R}^{T \times M}_+$.

Definition 1 The economy is said to allow for transfers if $a \in \mathcal{A}$ for any $a \in \mathbb{R}^{T \times M \times N}_+$ such that $\mathcal{P}(a) \in \mathcal{A}_X$,

A special case of an economy that allows for transfers is the *exchange economy*, in which aggregate production is fixed:

[4] Setting $T = 1$ yields the static model as a special case.

[5] The set $\mathbb{R}_{++}$ denotes the set of strictly positive real numbers, whereas $\mathbb{R}_+$ denotes the set of weakly positive real numbers.

[6] The measure-theoretic term *adapted* means that the process is well-defined with respect to the information available, at all points in time, see Duffie (2001), p. 21.

Definition 2 An *exchange economy* is an economy that allows for transfers, in which aggregate production is a singleton, $\mathcal{A}_X = \{X\}$.

Typically, a Walrasian market clearing mechanism is also specified for the exchange economy, leading to equilibrium prices and allocations. This mechanism is unimportant for our analysis and results, so there is no need for us to expand on this mechanism here.

Agents are expected utility maximizers. Agent n's expected utility under allocation a given probability vector q is

$$U^n(a|q^n) = \sum_{m=1}^{M} U_m^n(a) q_m^n, \qquad where \qquad U_m^n(a) = \sum_{t=1}^{T} u_{t,m}^n(a_{t,m,n}). \tag{1}$$

Here, we assume that each agent-, state-, and time-specific utility function, $u_{t,m}^n : \mathbb{R}_+ \to \mathbb{R}$, is strictly increasing, continuously differentiable, and strictly concave. Moreover, each agent's utility function, viewed as a process $u^n : \mathcal{T} \times \Omega \times \mathbb{R}_+ \to \mathbb{R}$ is required to be adapted.

With each allocation, we associate the utility matrix, $V = \mathcal{V}(a) \in \mathbb{R}^{M\times N}$, through the mapping $V_{m,n} = U_m^n(a)$, $1 \le m \le M$, $1 \le n \le N$, and define the *utility possibility set* (UPS), $\mathcal{U} = \mathcal{V}(\mathcal{A}) \subset \mathbb{R}^{M\times N}$. Since $\mathcal{A}$ is compact, so is $\mathcal{U}$. An important class of economies, as we shall see, are those with convex utility possibility sets.

Condition 1 The utility possibility set, $\mathcal{U}$, is convex.

The convexity condition for the utility possibility set is commonly imposed (see Mas-Colell et al. 1995). In an economy that allows for transfers, a sufficient condition for $\mathcal{U}$ to be convex is that the aggregate production set, $\mathcal{A}_X$ is convex.

Condition 2 The aggregate production set, $\mathcal{A}_X$, is convex.

The UPS-frontier, $\mathcal{G}$, consists of the elements in the utility possibility set that are not dominated, i.e., $g \in \mathcal{G}$, if there is no $h \in \mathcal{U}$, such that $h_{m,n} \ge g_{m,n}$ for all m and n with at least one inequality being strict. We will explicitly point out when additional assumptions are made about the aggregate production set and/or utility possibility set.

2.1 Agreement

We first study the case when there is agreement, $q^n = q$ for all n, as a benchmark case. A social planner has a Bergson welfare function over feasible allocations, $U(a|q, \lambda)$, defined by

$$U(a|q, \lambda) = \sum_{n=1}^{N} \lambda^n U^n(a|q), \tag{2}$$

where $\lambda \in S^N$ are Pareto weights in the planner's welfare function. Using the rules of matrix-vector multiplication, and denoting the transpose of the vector q by q', it then follows that

$$U(a|q,\lambda) = q'\mathcal{V}(a)\lambda, \tag{3}$$

i.e., that given an allocation, a, the welfare function is a bilinear mapping on the pair of probability vector and the vector of Pareto weights.

If $U(b|q,\lambda) > U(a|q,\lambda)$ for two allocations, $a, b \in \mathcal{A}$, we write $b >_q^\lambda a$, and if $U(b|q,\lambda) \geq U(a|q,\lambda)$, we write $b \geq_q^\lambda a$. With this notation, Pareto dominance and efficiency (see Mas-Colell et al. 1995) under agreement can be defined as follows:

Definition 3 (*Pareto dominance and efficiency*)

(i) Allocation b *Pareto dominates* a given q, $b \succ_q a$, if $b >_q^\lambda a$ for all $\lambda \in S^N$.
(ii) Allocation b is not Pareto dominated by a given q, $b \succeq_q a$, if $b \geq_q^\lambda a$ for some $\lambda \in S^N$.
(iii) Allocation b is *Pareto efficient* given q, if $\forall a \in \mathcal{A} : b \succeq_q a$.

We denote by E_q the set of all Pareto efficient allocations given probability vector q. Define

$$M_{q,\lambda} = \arg\max_{a\in\mathcal{A}} q'\mathcal{V}(a)\lambda. \tag{4}$$

Since $\mathcal{A}$ is compact, $\mathcal{V}(\mathcal{A})$ is also compact and $M_{q,\lambda}$ is nonempty. Also, $M_{q,\lambda} \subset E_q$ for any $\lambda \in S^N$, and thus E_q is nonempty.

An equivalent definition of $b \succ_q a$ is that $b \geq_q^\lambda a$ for all $\lambda \in \bar{S}^N$, with the inequality being strict for at least one such λ. Also, an equivalent definition of allocation b being Pareto inefficient given q is that

$$\exists a \in \mathcal{A}, \forall \lambda \in S^N : a >_q^\lambda b. \tag{5}$$

2.2 Disagreement

Under disagreement, agent n's belief is denoted by q^n. In line with assumptions made in Brunnermeier et al. (2014) and Heyerdahl-Larsen and Walden (2022), a nonempty closed set, $\mathcal{Q}_R \subset S^M$ of beliefs is viewed as "reasonable." Although this *reasonable belief set* plays a very similar similar role in the definition in the two papers, the interpretation of what it represents is somewhat different. The measures in Brunnermeier et al. (2014) (belief neutral efficiency and belief neutral, inefficiency, discussed later in this paper) are natural if the social planners can determine a $q \in \mathcal{Q}_R$ that is the correct belief.

The measure in Heyerdahl-Larsen and Walden (2022) is natural if the social planner cannot identify such a q, but instead views each $q \in \mathcal{Q}_R$ as possible. The efficiency concepts we study all take into account the welfare associated with all reasonable beliefs, in line with Brunnermeier et al. (2014), Gilboa et al. (2014), Gayer et al. (2014), and Heyerdahl-Larsen and Walden (2022). The special case for which $q^n = q$ for all n reduces to the agreement economy, in which case we assume that the planner's reasonable beliefs set is $\mathcal{Q}_R = \{q\}$.

We are agnostic about the choice of $\mathcal{Q}_R$ in the general case, and make as few assumptions as possible about this set. Brunnermeier et al. (2014) suggests using the convex hull of agents' individual beliefs. Specifically, for a set $X \subset \mathbb{R}^K$, define $CH(X) = \left\{ \sum_{k=1}^{K} \rho_k x_k \, : \, \rho \in \bar{S}^K, x_k \in X, 1 \le k \le K \right\}$. The set suggested in Brunnermeier et al. (2014) is then

$$\mathcal{Q}_R^{CH} = CH\left(\left\{q^1, q^n, \dots, q^N\right\}\right). \tag{6}$$

This choice has several appealing properties, although other choices may be more appropriate in some situations. For example, the set of beliefs, $\mathcal{Q}_R = \{q^1, \dots, q^N\}$, is an alternative candidate. Both of these choices satisfy the following condition:

Condition 3 For all n, $q^n \in \mathcal{Q}_R$.

We provide another example in Sect. 4 where a case could be made that neither the convex hull of beliefs or the set of beliefs is the right choice of $\mathcal{Q}_R$. The only restriction on $\mathcal{Q}_R$ we impose is that the set contains a single element, $\mathcal{Q}_R = \{q\}$ if and only if agents agree, and $q^n = q$ for all $1 \le n \le N$. Thus, under disagreement, $|\mathcal{Q}_R| \ge 2$. Later, we will impose restrictions on $\mathcal{Q}_R$, which we will then explicitly point out.

We denote the efficiency concept used in Heyerdahl-Larsen and Walden (2022) *Incomplete Knowledge efficiency*, or simply *IK-efficiency*, since it is based on the assumption that the social planner has incomplete knowledge about the true probabilities in the economy.

Following Heyerdahl-Larsen and Walden (2022):

Definition 4 (*IK-dominance*) Allocation *b IK-dominates a with respect to Pareto weights* λ, $b \succ^\lambda a$, if:

$$(\forall q \in \mathcal{Q}_R \, : \, b \ge_q^\lambda a) \text{ and } (\exists q \in \mathcal{Q}_R \, : \, b >_q^\lambda a),$$

From Definition 4 it follows that an allocation dominates another under this concept if, given Pareto weights, it is never strictly dominated under any reasonable probability, and there exist a probability measure under which it strictly dominates the other allocation. We also define weak IK-dominance, $a \succeq^\lambda b$, if $\neg(b \succ^\lambda a)$. Here, "$\neg$" is the logical negation symbol.

IK-efficiency is now defined as follows:

Definition 5 (*IK-efficiency*) Allocation a is *IK-inefficient* if $\forall \lambda \in S^N, \exists b \in \mathcal{A}$: $b \succ^\lambda a$. Equivalently, a is IK-inefficient if

$$\forall \lambda \in S^N, \exists b \in \mathcal{A}, \forall q \in \mathcal{Q}_R : b \geq_q^\lambda a, \tag{7}$$

where the inequality is strict for at least one q.

An allocation that is not IK-inefficient is called IK-efficient. We denote the set of IK-efficient allocations by *IKE*, and the set of IK-inefficient allocations is then $IKE^c = \mathcal{A} \setminus IKE$. An IK-inefficient allocation is thus one for which whatever are the Pareto weights in the welfare function, there exists another allocation that is not dominated by the first regardless of q in the set of reasonable beliefs, and that dominates the first for some reasonable q.

It is straightforward to verify that an equivalent definition for allocation a to be IK-efficient is that

$$\exists \lambda \in S^N, \forall b \in \mathcal{A} : a \geq^\lambda b,$$

i.e., that

$$\exists \lambda \in S^N, \forall b \in \mathcal{A} : (\exists q \in \mathcal{Q}_R : a >_q^\lambda b, \text{ or } \forall q \in \mathcal{Q}_R : a \geq_q^\lambda b). \tag{8}$$

An IK-inefficient allocation is operational for a planner who is not able to take a stand on which q is correct among the set of reasonable beliefs, in that whatever the planner's Pareto weights are there is another allocation b that improves upon a without taking a stand on q. The incomplete knowledge concept thus requires the social planner to have a well-defined λ, but does not require that (s)he takes a stand on a unique q among the set of reasonable beliefs, $\mathcal{Q}_R$.

Next, we introduce two concepts of belief neutral efficiency, based on Brunnermeier et al. (2014). In our economic environment, belief neutral efficiency and inefficiency are defined as follows:

Definition 6 (*Belief neutral efficiency*)

- Allocation a is *belief neutrally inefficient*, $a \in BNI$, if $\forall q \in \mathcal{Q}_R, \exists b \in \mathcal{A} : b >_q a$, i.e., if

$$\forall q \in \mathcal{Q}_R, \exists b \in \mathcal{A}, \forall \lambda \in S^N : b >_q^\lambda a. \tag{9}$$

- Allocation a is *belief neutrally efficient*, $a \in BNE$, if $\forall q \in \mathcal{Q}_R, \forall b \in \mathcal{A} : a \geq_q b$, i.e., if

$$\forall q \in \mathcal{Q}_R, \forall b \in \mathcal{A}, \exists \lambda \in S^N : a \geq_q^\lambda b. \tag{10}$$

In words, an allocation, a, is belief neutrally inefficient if for every reasonable belief, q, there is another allocation, b, that is strictly better regardless of the Pareto weights, λ. The set of belief neutrally inefficient allocations is in general a strict subset of the complement of the set of belief neutral efficient allocations, i.e., $BNI \subsetneq BNE^c$. There may thus be allocations that are neither belief neutrally efficient, nor belief neutrally inefficient. To avoid the cumbersome terminology of allocations being "not belief neutrally inefficient," we denote such allocations *weakly belief neutrally efficient*:

Definition 7 Allocation a is *weakly belief neutrally efficient*, $a \in WBNE$, if $a \notin BNI$, i.e., if

$$\exists q \in \mathcal{Q}_R, \forall b \in \mathcal{A}, \exists \lambda \in S^N \,:\, a \geq_q^\lambda b. \tag{11}$$

We stress that the weakly belief neutral terminology is not used in Brunnermeier et al. (2014).

It immediately follows that equivalent definitions of $WBNE$ and BNE are:

$$WBNE = \cup_{q \in \mathcal{Q}_R} \; E_q, \tag{12}$$

$$BNE = \cap_{q \in \mathcal{Q}_R} \; E_q, \tag{13}$$

and therefore that $BNE \subset WBNE$.

Note that the roles of the Pareto weights, λ, and probabilities, q, are in some sense dual in the definitions of IK-efficiency and belief neutral efficiency. Specifically, under the IK efficiency concept, the alternative allocation is allowed to vary with λ but not q, whereas under belief neutral efficiency it is allowed to vary with q but not λ.

Next, we define a concept that is related to the unanimity Pareto dominance concept discussed in Gayer et al. (2014):

Definition 8 (*U-efficiency*)

Allocation a is *U-inefficient* if $\exists b \in \mathcal{A}, \forall q \in \mathcal{Q}_R \,:\, b >_q a$, i.e., if

$$\exists b \in \mathcal{A}, \forall q \in \mathcal{Q}_R, \forall \lambda \in S^N \,:\, b >_q^\lambda a. \tag{14}$$

An allocation, a, that is not U-inefficient is called U-efficient, i.e.,

$$\forall b \in \mathcal{A}, \exists q \in \mathcal{Q}_R, \exists \lambda \in S^N \,:\, a \geq_q^\lambda b. \tag{15}$$

We denote the set of U-efficient allocations by UE. U-inefficiency is thus a strong form of inefficiency, since it requires the existence of a unique allocation that dominates a current allocation, regardless of both Pareto weights, $\lambda \in \mathcal{S}^N$, and probabilities, $q \in \mathcal{Q}_R$. This is in contrast to IK-inefficiency and belief neutral inefficiency, which both allow the alternative allocation to vary with one of these parameters. It follows that $WBNE \subset UE$, and $IKE \subset UE$.

Although U-efficiency has some similarity with the Unanimity Pareto concept introduced in Gayer et al. (2014), there are also differences: Gayer et al. (2014) focus on agents involved in a transaction, and require all those agents to be *strictly* better off. More importantly, theirs is a dominance concept that can be used an allocation with another, rather than an efficiency measure. Finally, their concept requires each agent to be better off given his or her *own* beliefs, for a reallocation to be identified

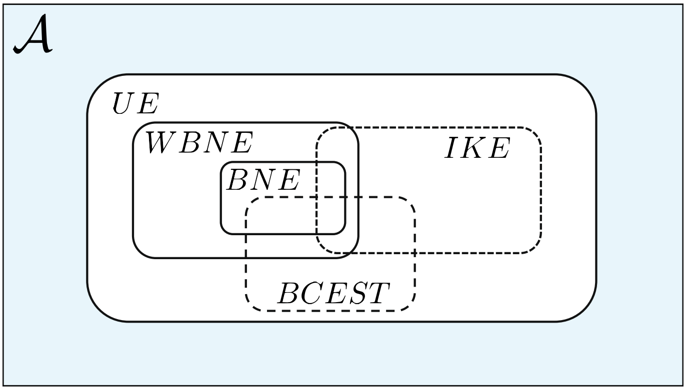

Fig. 1 Venn diagram of relationship between the different concepts in the general economy. For visual clarity, the lines for the *IKE* and *BCEST* sets are dashed. These relations hold in the general economy, without further restrictions on the economic environment

as an improvement, and is therefore somewhat different in spirit from the other measures, which all are based on using a common belief for all agents, among a set of reasonable beliefs. This final condition is also imposed in the definition of a related dominance concept discussed in Gayer et al. (2014), and also in Gilboa et al. (2014), namely so-called No-Betting Pareto Dominance. Gayer et al. (2014) discuss how their Pareto dominance concepts are related to belief neutral efficiency. Our focus is mainly on comparing efficiency measures.

We also introduce the set *BCEST*, following the analysis in Blume et al. (2018) and named after the authors' initials, of maximal solutions with respect to the Bergson welfare function.[7]

Definition 9 The set *BCEST* is defined as:

$$BCEST = \cup_{\lambda \in S^N} \arg\max_{a \in \mathcal{A}} \mathcal{W}(a, \lambda), \tag{16}$$

where

$$\mathcal{W}(a, \lambda) = \min_{q \in \mathcal{Q}_R} U(a|q, \lambda). \tag{17}$$

This set contains all the outcomes a planner with the welfare function in Blume et al. (2018) may view as optimal, depending on Pareto weights λ, given that all allocations $a \in \mathcal{A}$ are implementable. We note that this measure is conservative with respect to the set of reasonable beliefs, being defined over the minimum over all such beliefs. It is easy to show that $BCEST \subset UE$.[8]

[7] We focus on the Bergson welfare formulation, while noting that Blume et al. (2018) also allow for more general aggregators.

[8] For $a \in BCEST$, there exists $\lambda \in S^N$, such that for all $b \in \mathcal{A}$, $\min_{q \in \mathcal{Q}_R} U(b|q, \lambda) \leq \min_{q \in \mathcal{Q}_R} U(a|q, \lambda)$. Because of compactness, these minima are realized, so for some $q^* \in \mathcal{Q}_R$, $\min_{q \in \mathcal{Q}_R} U(b|q, \lambda) = U(b|q^*, \lambda)$

The relationships between the different efficiency sets in the general economy are summarized in Fig. 1.

Finally, we introduce the concept of an *Arrow optimum*, a term used in Starr (1973). As discussed in that paper, an allocation is an Arrow optimum if it is Pareto efficient with respect to the expected utilities of agents, based on their own individual beliefs.[9] In our setting, the set of Arrow optima are identified by replacing the welfare function (2) by

$$U(a|\lambda, \mathbf{q}) = \sum_{n=1}^{N} \lambda^n U^n(a|q^n), \tag{18}$$

where $\mathbf{q} = (q^1, q^2, \ldots, q^N) \in \prod_{n=1}^{N} S^M$ represents all agents' beliefs. If $U(b|\lambda, \mathbf{q}) > U(a|\lambda, \mathbf{q})$ for two allocations, we then write $b >_{\mathbf{q}}^{\lambda} a$, and if $U(b|\lambda, \mathbf{q}) \geq U(a|\lambda, \mathbf{q})$, we write $b \geq_{\mathbf{q}}^{\lambda} a$. The definition of Arrow dominance and Arrow optimum are then identical as in Definition 3 (i.e., holding for all $\lambda \in S^N$), but with the single probability vector q replaced by the N-tuple of probability vectors, $\mathbf{q}$, with each agent using his/her own belief. The set of Arrow optima, given beliefs $\mathbf{q}$, is denoted by $E_{\mathbf{q}}^{A}$. A similar argument as for E_q implies that $E_{\mathbf{q}}^{A}$ is nonempty.

Intuitively, Arrow optima, being based on agents' individual beliefs, allow for speculative outcomes in which there is significant variation in allocations across states because of disagreement. This occurs because agents dismiss consumption in states they subjectively believe are very unlikely. For such speculative allocations, it is typically objectively known that many agents will end up with low consumption, since all agents' beliefs cannot be correct, as discussed in Brunnermeier et al. 2014 and Gilboa et al. 2014.[10] The efficiency measures introduced here—IK-efficiency, belief neutral efficiency (strong and weak), and U-efficiency measures—are all designed to rule out such allocations as being inefficient, by forcing the same probability measure to be used across agents when comparing allocations, although the measures differ on the specific details.

We have not included the set of Arrow optima in Fig. 1. This set may be quite different from the other efficiency sets, in general being neither a subset or a superset of any of the other sets, as shown in Proposition 2 and the following discussion, below.

Footnote 8 (continued)

$\leq U(a|q^*, \lambda)$. Equivalently, $\forall b \in \mathcal{A}, \exists q^* \in \mathcal{Q}_R : a \geq_{q^*}^{\lambda} b$, and therefore one sees from (14) that a is not UK-inefficient. So, $BCEST \subset UE$.

[9] Harris (1978) calls such an allocation *ex ante* efficient.

[10] For example, two risk-averse agents who have drastically different beliefs about the probability for heads being the outcome of a coin toss may both prefer an allocation where, depending on the outcome of the coin toss, one agents gets all of the consumption good and the other agent starves, over one in which they share the good equally in both states.

3 Comparing the different measures

Focusing first on the agreement economy, we have:

Proposition 1 *In the agreement economy with probability vector q*:

(i) *In general*, $IKE = BCEST \subset E_q = E^A_{\mathbf{q}} = WBNE = BNE = UE$.
(ii) *If the economy allows for transfers and the utility possibility set is convex*, *then* $IKE = BCEST = E^A_q = E^A_{\mathbf{q}} = WBNE = BNE = UE$.

Thus, in economies that allow for transfers with convex utility possibility sets, the efficiency concepts are all identical in the agreement economy, whereas *IKE* and *BCEST*, may be strict subsets of the other efficiency sets otherwise.

We move on to the disagreement economy, in which $q^n \neq q^{n'}$ for at least two agents, and consequently $|\mathcal{Q}_R| \geq 2$. Our first result shows that in general it is not be possible for U-efficiency to rule out all speculative allocations that are Arrow optima, whereas the other efficiency concepts typically *do* rule out such speculative allocations:

Proposition 2 *In the disagreement economy*:

(i) *If Condition 3 is satisfied*, *i..e.*, $q^n \in \mathcal{Q}_R$ *for all agents*, *then* $E^A_{\mathbf{q}} \subset UE$.
(ii) *Any Arrow optimum*, $a \in E^A_{\mathbf{q}}$, *in an economy that allows for transfers*, *in which two agents that disagree about the relative likelihood of two states to occur are allocated strictly positive amounts of the consumption goods in both those states at some time*, *is neither IK-efficient*, *nor weakly belief neutral efficient*, *i.e.*, $a \notin IKE \cup WBNE$.

In economies in which the first welfare theorem holds, competitive equilibria are Arrow optima. The proposition implies that such speculative equilibria may be identified as U-efficient. Hence, whereas both IK-efficiency and weak belief neutral efficiency are designed to rule out all speculative allocations, the U-efficiency criterion is not sufficiently strong to do so.[11] Part (iii) of the proposition shows that both *IKE* and *WBNE* rule out all such speculative allocations.

A simple counterexample, where $E^A_{\mathbf{q}} \not\subset UE$, is the following: There are two allocations, $\mathcal{A} = \{a, b\}$, two agents, and two states. Agent beliefs are defined by $q^1_1 = q^2_2 = 0.9$. Allocation a is risk-free, $U^1(a|q) = U^2(a|q) = 1$, for all q. Agent utilities under allocation b are $U^1(b|q) = 1.5q_1$, $U^2(b|q) = 1.5q_2$. The reasonable belief set is $\mathcal{Q}_R = \{(q_1, 1-q_1)' : 0.4 \leq q_1 \leq 0.6\}$.

[11] Belief neutral inefficiency of Arrow optima—part (iii) of Proposition 2—actually follows from the analysis in Starr (1973), see Starr's Corollary 3.1 on page 81.

Now, $\lambda U^1(a|q) + (1-\lambda)U^2(a|q) = 1$ regardless of q, whereas $\sup_{\lambda \in S^2} \max_{q \in \mathcal{Q}_R} \lambda U^1(b|q) + (1-\lambda)U^2(b|q) = 1.5 \times 0.6 = 0.9$. So regardless of Pareto weights and reasonable beliefs, planner welfare is higher under a than b, and therefore $UE = \{a\}$. However, regardless of λ, $U(b|\lambda, \mathbf{q}) = 1.35 > U(a|\lambda\mathbf{q}) = 1$, so $E^A_{\mathbf{q}} = \{b\}$. Thus, UE and $E^A_{\mathbf{q}}$ are disjoint in this example.

It also follows almost immediately that in the environment of Gayer et al. (2014), allocations that belong to $E^A_{\mathbf{q}}$, or to UE as long as Condition 3 is satisfied, can not be Unanimity Pareto dominated by any reallocation, implying that the set of such Unanimity Pareto undominated allocations is large.[12] Similarly, allocations that belong to $E^A_{\mathbf{q}}$ cannot be No-Betting Pareto dominated by any reallocation.

We next relate *WBNE* and *IKE*. In general, these efficiency concepts differ (see examples in Sect. 4), but under additional conditions, that are satisfied in several work-horse models, they coincide. We introduce the following conditions:

Condition 4 The set of reasonable beliefs, $\mathcal{Q}_R$, is convex.

Condition 5 Strict dominance condition: For any $q \in \mathcal{Q}_R$ and $\lambda \in S^N$, $M_{q,\lambda}$ defined in (4) is a singleton, $|M_{q,\lambda}| = 1$.

The convexity condition for the set of reasonable beliefs is satisfied under the assumptions made in Brunnermeier et al. (2014), see Eq. (6). The strict dominance condition states that each allocation that maximizes $q'\mathcal{V}(a)\lambda$ for a reasonable probability and associated vector, strictly dominates all other allocations, w.r.t. q and λ. A sufficient condition for strict dominance is that the production set is convex and the planner's problem strictly is concave, as shown in the following proposition:

Proposition 3 *In an economy that allows for transfers, in which the aggregate production set is convex, the strict dominance condition is satisfied.*

We note that one way of ensuring convexity of the utility possibility set is by allowing for randomization, see Yaari (1981). Specifically, if the planner uses a randomization device to choose between allocations $a_1, \dots, a_K$ with probabilities $\rho_1, \dots, \rho_K$, the associated utility matrix is $\sum_{k=1}^K \rho_k \mathcal{V}(a_k)$, which consequently belongs to $\mathcal{U}$. This rationale for a convex utility possibility is of course more subtle in a setting with disagreement than in Yaari (1981), because agents need to objectively agree about the probabilities of the randomization device to ensure that new disagreement is not introduced in the process. If, however, such *objective randomization* is possible, convexity of the utility possibility set follows.

The following two results relate *IKE* and *WBNE*:

12 For allocation b to Unanimity Pareto dominate a in their environment, the allocation must both Arrow dominate and dominate with respect to U-efficiency under the reasonable belief set $\mathcal{Q}_R = \{q^1, \dots, q^N\}$, according to their definition.

Proposition 4 *In an economy that allows for transfers, in which the utility possibility set and reasonable belief set are both convex,* $IKE \subset WBNE$.

Proposition 5 *In an economy that allows for transfers, in which the utility possibility set is convex, and in which the strict dominance condition is satisfied,* $WBNE \subset IKE$.

Propositions 4 and 5 together show that in economies with transfers that satisfy Conditions 1, 4, and 5, IK-inefficiency and belief neutral inefficiency are equivalent. Several standard work-horse models fall within this class of economies. For instance, the early production based asset pricing models in Brock (1979), Mehra and Prescott (1980) and Cox et al. (1985) have similar setups and form the basis for many of the general equilibrium models in the field of macro-finance. It also follows from (12) that both these efficiency sets are nonempty when these conditions are satisfied.

As suggested by (12,13), belief neutral efficiency puts substantially stronger restrictions on allocations than weak belief neutral efficiency, and we may therefore expect *BNE* to be a "small" set in many cases. Indeed, a belief neutral efficient allocation must be efficient for *all* reasonable q, as seen from (13). In economies for which the optimal aggregate production depends on q, *BNE* may therefore be empty. We have

Proposition 6 *In an economy that allows for transfers and has a convex production set (i.e., satisfies Condition 2), BNE is nonempty if and only if* $\cap_{q\in\mathcal{Q}_R}\mathcal{P}(E_q)$ *is nonempty.*

Proposition 6 shows that what determines whether *BNE* is nonempty in the production economy with transfers is whether all agents can agree on what is an efficient aggregate level of production. The necessity of the condition is trivial, and sufficiency follows from convexity of the production set, as shown in the proof of the proposition.

A straightforward example where Proposition 6 applies is the exchange economy, in which the production set is a singleton, and belief neutral efficiency and inefficiency, IK-efficiency, and the set *BCEST* are all equivalent. We have:

Proposition 7 *In the exchange economy,*

(i) $E_q = E_{q'}$ *for all* $q, q' \in S^M$,
(ii) $IKE = BCEST = WBNE = BNE$.

It is reassuring for the use of multiple-beliefs based efficiency measures that the technical differences in how the measures are defined are unimportant in the exchange economy. However, as our previous results suggest, it would be incorrect to conclude that they do not matter in more general economic environments, that include production. In the next section, we provide several examples of such settings, where the choice of measure *does* matter. Note also that in light of Proposition 2(ii), Proposition 7 implies that Arrow optima that contain speculation are not in any of the other efficiency sets, in the exchange economy.

We end this section with an example that shows how the exchange economy and production economy differ.

3.1 An example

There is one date, $M = 2$ states and $N = 2$ agents, both of whom have square root utility, $u^n(x) = \sqrt{x}$, $n = 1, 2$. Agent 1 believes the probability for state 1 to occur is $q_1^1 = p^1$, and agent 2 believes it us $q_1^2 = p^2$. Transfers are allowed, and initially we focus on the case where the aggregate production set is a singleton, i.e., on the exchange economy setting, There is half a unit in total of the consumption good in both states, so $\mathcal{A}_X = \{(1/2, 1/2)\}$. It follows that the social planner's welfare function (2) in this case may be written as

$$U(\lambda, q) = \sqrt{\frac{1}{2}}\Big(\lambda\sqrt{z_1}q_1 + (1-\lambda)\sqrt{1-z_1}q_1 + \lambda\sqrt{z_2}(1-q_1) + (1-\lambda)\sqrt{1-z_2}(1-q_1)\Big). \quad (19)$$

Here, $q = (q_1, q_2)$ is the common probability vector used for both agents, $\lambda_1 = \lambda$ represents the planner's weight given to agent 1, and $\lambda_2 = 1 - \lambda$ the weight given to agent 2. Moreover, z_1 and z_2 define the shares of total consumption allocated to agent 1 in states 1 and 2, respectively, corresponding to the allocation $a = (a_{1,1}, a_{2,1}, a_{1,2}, a_{2,2}) = a(z_1, z_2) \overset{\text{def}}{=} \frac{1}{2}(z_1, z_2, 1 - z_1, 1 - z_2)$, where $a_{m,n}$ represents agent n's allocation in state m. So, z_1 and z_2 represent agent 1's consumption share in state 1 and 2, respectively. The set of feasible allocations is then $\mathcal{A} = \{a(z_1, z_2) : 0 \le z_1 \le 1, 0 \le z_2 \le 1\}$.

Given belief q and weight λ, one verifies that this leads to the following optimal consumption shares

$$z_1 = z_2 = Z(\lambda), \qquad \text{where },$$
$$Z(\lambda) \overset{\text{def}}{=} \frac{1}{1 + \left(\frac{1}{\lambda} - 1\right)^2}, \quad (20)$$

i.e., those shares that give each agent the same share across the two states.

One also verifies that the function Z is a strictly increasing bijection on the open unit interval, which implies that the efficient allocations are

$$E_q = E \overset{\text{def}}{=} \{a(x, x) : x \in (0, 1)\}.$$

So, in line with Proposition 7(a), E_q does not depend on beliefs. This is because the optimization of (19) is effectively carried out state-by-state, regardless of what the probabilities for those states are, leading to identical allocation rules in each state. From (12, 13), it moreover follows that $WBNE = BNE = E$, and from Propositions 4 and 5 that $IKE = WBNE$, since Conditions 1–5 are satisfied. Finally, from (16,17) it follows that $BCEST = E$, because, as elaborated upon in the proof of Proposition 7, the allocations that maximize $\mathcal{W}$ are exactly those in E, regardless of which q minimizes $U(a|q,\lambda)$ in (17). These results are all in line with Proposition 7(b), showing that in the exchange economy restricting agent beliefs to be common in the planner's problem resolves the issue of speculation, and that the specific assumptions for how to treat the multiple common beliefs in the reasonable belief set, where the efficiency measures differ, are irrelevant in this case.

Note that the set E is different from the set of Arrow optima $E^A_{\mathbf{q}}$, as defined in (18), when there is disagreement, $p^1 \neq p^2$. Specifically, the set of Arrow optima is defined by maximizing:

$$U = \sqrt{\frac{1}{2}}\Big(\lambda\sqrt{z_1}p^1 + (1-\lambda)\sqrt{1-z_1}p^2 + \lambda\sqrt{z_2}(1-p^1) \\ +(1-\lambda)\sqrt{z_2}(1-p^2)\Big), \qquad \lambda \in (0,1),$$

leading to the optimal consumption shares:

$$z_1 = \frac{1}{1+\left(\frac{1}{\lambda}-1\right)^2\left(\frac{p^2}{p^1}\right)^2},$$
$$z_2 = \frac{1}{1+\left(\frac{1}{\lambda}-1\right)^2\left(\frac{1-p^2}{1-p^1}\right)^2}.$$

It is easily seen that $z_1 \neq z_2$ when $p^1 \neq p^2$, so $E^A_{\mathbf{q}}$ and E are completely disjoint when there is disagreement. The difference between the sets is that $E^A_{\mathbf{q}}$, being based on the two agents' individual beliefs, leads to speculative reallocation so that agents consume more in the states they are relative optimistic about compared with the allocations identified in E, which rule out such speculation. In other words, all the efficiency measures succeed in ruling out the speculative Arrow optima and identify the nonspeculative set E.

We next study a variation of this example, which includes production. Specifically, the set-up is as before, but now with two linear production technologies: the first generates a unit of consumption (only) in state 1, per unit of capital invested, whereas the second generates a unit of consumption (only) in state 2. One unit of capital is available. If k is invested in the first technology and the remaining $1-k$ in the second, together with the consumptions shares z_1 and z_2 in states 1 and 2, respectively, this then implies the following allocation $a(k,z_1,z_2) \stackrel{\text{def}}{=} (kz_1,(1-k)z_2,k(1-z_1),(1-k)(1-z_2))$. The set of feasible allocations in this production economy is

$$\mathcal{A} = \{a(k,z_1,z_2) : 0 \le k \le 1, 0 \le z_1 \le 1, 0 \le z_2 \le 1\}.$$

Note that the exchange economy above corresponds to the special case for which $k = \frac{1}{2}$ is fixed.

To limit the number of parameters, we focus on the case when $p^1 = 1-p^2 = p$, so that agent beliefs are symmetric in that agent 1's belief about one state is the same as agent 2's belief about the other. Without loss of generality, we assume that $p \leq \frac{1}{2}$. The social planner's welfare can now be written as

$$\begin{aligned}U(\lambda, q) =& \lambda q_1 \sqrt{kz_1} + (1-\lambda) q_1 \sqrt{k(1-z_1)} \\ &+ \lambda(1-q_1)\sqrt{(1-k)z_2} + (1-\lambda)(1-q_1)\sqrt{(1-k)(1-z_2)} \\ =& q_1\sqrt{k}\Big(\lambda\sqrt{z_1} + (1-\lambda)\sqrt{1-z_1}\Big) + (1-q_1)\sqrt{1-k}\Big(\lambda\sqrt{z_2} + (1-\lambda)\sqrt{1-z_2}\Big).\end{aligned}$$

One verifies that the optimal shares are the same as in the exchange economy, $z_1 = z_2 = Z(\lambda)$, as defined in (20), independently of k and q_1, and that the optimal investment k is $k = Z(q_1)$, regardless of λ. As a consequence,

$$E_q = \{a(Z(q_1), x, x) \,:\, x \in (0,1)\}.$$

Now, given that agent beliefs are $q^1 = (p, 1-p)$ and $q^2 = (1-p, p)$, and choosing the reasonable belief set to be the convex hull of individual beliefs, it follows that $\mathcal{Q}_R = \{(q, 1-q) \,:\, q \in [p, 1-p]\}$.

There is a major difference compared with the exchange economy here, in that beliefs matter for what is optimal aggregate outcomes. When q^1 is higher, it is optimal to invest more in the first production technology, which produces in state 1. As a consequence, E_q varies with q in this production economy. In other words, in addition to ruling out speculation, which is done in the same way regardless of q, E_q takes a stand on how to optimally allocate productive resources in this setting, which depends on q.

Since the function Z is strictly increasing, it follows immediately that $E_q \cap E_{q'} = \emptyset$ when $q \neq q'$, and therefore, via (13), that $BNE = \emptyset$. Indeed, it is clear from this example that BNE is empty because $\cap_{q \in \mathcal{Q}_R} \mathcal{P}(E_q) = \emptyset$, in line with Proposition 6. In this environment, BNE is therefore of limited use.

It is straightforward to verify that Conditions 1–5 are satisfied, so Propositions 4 and 5 together with (12) imply

$$WBNE = IKE = \cup_{q \in \mathcal{Q}_R} E_q.$$

Finally, note from (17) that

$$\begin{aligned}\mathcal{W}(a(k,x,x), \lambda) =& \min_{q \in \mathcal{Q}_R} (q_1\sqrt{k} + (1-q_1)\sqrt{1-k})(\lambda\sqrt{x} + (1-\lambda)\sqrt{1-x}) \\ =& \frac{1}{2}(\sqrt{k} + \sqrt{1-k})(\lambda\sqrt{x} + (1-\lambda)\sqrt{1-x})),\end{aligned}$$

which is realized for $q = (1/2, 1/2)$. This, in turn, implies that $\arg\max \mathcal{W}$ is realized for $k = \frac{1}{2}$, and thus via (16) that

$$BCEST = E_{(1/2,1/2)} = \{a(1/2, x, x) \,:\, x \in (0,1)\}.$$

So, $BCEST$ is a strict subset of $WBNE$ and IKE here, by forcing investments $k = 1/2$. In contrast, $WBNE$ and IKE view all investments $k \in [Z(p^1), Z(1-p^1)]$ as efficient.

The planner's approach when defined by $BCEST$ is to look at the worst-case probability scenario among the set of reasonable beliefs, as shows in the minimization used to define $\mathcal{W}$. This functional form has a similarity with the max–min ambiguity aversion model of Gilboa and Schmeidler (1989), and it narrows down the set of efficient allocations. Indeed, the planner ends up taking a strong stand on which belief in $\mathcal{Q}_R$ to use. In contrast, for *WBNE* and *IKE* the planner is unwilling to select a specific belief, viewing any allocation that is efficient under *some* reasonable belief as efficient. Which of these approaches is most appropriate will depend on the circumstances.

4 Further discussion and examples

To further highlight the properties of these efficiency measures, and their differences, we study several examples. We first study an intertemporal (two date) production economy with many states, that may be viewed as a work-horse model for which $WBNE = IKE$. We then move on to other examples, outside of the class of economies that satisfy Conditions 1–5. These examples, in line with Fig. 1, show how the subtle differences between how the efficiency measures are defined become more important when outside of the class of economies that satisfy Conditions 1–5.

4.1 A general production economy in which *WBNE* = *IKE*

In Appendix C, we describe a general work-horse production economy that allows for transfers, which is a variation of the economy studied in Heyerdahl-Larsen and Walden (2022). Both the reasonable belief set and the aggregate production set are convex, and it therefore follows from

We consider a production economy with two dates, $t = 1, 2$, that allows for transfers. There are M possible states, and $N > 1$ agents, who disagree. The state is revealed at $t = 2$, so we require that $a_{m,n,1} = a_{m',n,1}$ for $1 \le m, m' \le M$, for all $a \in \mathcal{A}$. Moreover, we assume that agents have strictly concave utility, i.e., that the functions $u^n_{m,t}$ are strictly concave for all n and $m, t \in \{1, 2\}$, and that Condition 4 is satisfied, i.e., that the set of reasonable beliefs is convex.

From Propositions 3–5, it follows that IK-efficiency and weak belief neutral efficiency coincide in this economy,

$$IKE = WBNE.$$

When agents have power utility (constant relative risk aversion, with risk aversion coefficient not equal to one), *BNE* is typically empty,

$$BNE = \emptyset.$$

Moreover, the conditions of Proposition 2 (i) and (ii) are satisfied, so all Arrow optima are U-efficient, as well as IK-inefficient and weakly belief neutrally inefficient.

Table 1 Four allocations in economy with two agents and two states

Allocation,	a_1		a_2		a_3		a_4	
agent	1	2	1	2	1	2	1	2
State 1	1	1	9	9	0	0	1.9	0.1
State 2	1	1	0	0	9	9	0.1	1.9
Ex ante utility	1	1	2.7	0.3	0.3	2.7	1.27	1.27

Allocation a_1 and a_4 are based on investments in the risk-free technology, with equal (a_1) and unequal (a_4) sharing between agents in the two states. Allocation a_2 and a_3 both have equal sharing, but invest in risky technologies that pay off in state 1 and 2, respectively

4.2 A production economy with non-convex utility possibility set

There are two risk averse expected utility maximizing agents, three mutually exclusive production technologies, one date, and two states. The expected utility of agent $n \in \{1, 2\}$ is

$$U^n = q\sqrt{c_1^n} + (1-q)\sqrt{c_2^n}, \tag{21}$$

where c_m^n is the consumption of agent n in state $m \in \{1, 2\}$, and q is the probability for state 1. The first production technology is risk free and generates a total output of 2 units of the consumption good in either state. The second technology is risky, generating 18 units in state 1 and 0 units in state 2, as is the third technology which generates 18 units in state 2 and 0 in state 1. The economy neither permits transfers, nor objective randomization over events.

There are 4 possible allocations, captured by the set $\mathcal{A} = \{a_1, a_2, a_3, a_4\}$. In allocations a_i, $i = 1, 2, 3$, the consumption good is divided equally between the two agents, after choosing production technology i. In allocation a_4, the risk-free production technology is used, as in allocation a_1, but agent 1 receives 1.9 in state 1 and agent two receives 0.1, whereas agent 2 receives 1.9 in state 2 and agent 1 receives 0.1. The consumption by the two agents for different allocations and states is shown in Table 1. Allocation a_4 thus allows for speculation.

The two agents disagree about the probability for state 1 to occur, q. Agent 1 believes that the probability is $q^1 = 0.9$, whereas agent 2 believes it is $q^2 = 0.1$. The planner, not knowing which beliefs are correct, views any probability in the interval $\mathcal{Q}_R = [0.1, 0.9]$ as reasonable.

It is easy to verify that the only allocation that is not an Arrow optimum is a_1, which both agents agree is dominated by a_4, based on their different beliefs. Of course, both agents also agree that the welfare improvement is speculative, and that whatever the true q is, any allocation in which individual consumption shares vary across states can be improved upon by risk sharing. Thus, a_4 is inefficient whenever the same q is used for both agents. This is the speculative inefficiency that is captured by the novel efficiency measures.

In Fig. 2, the right panel compares allocation a_1 with a_4. The horizontal (black) line represents the (same) utility of the two agents under the risk free allocation a_1,

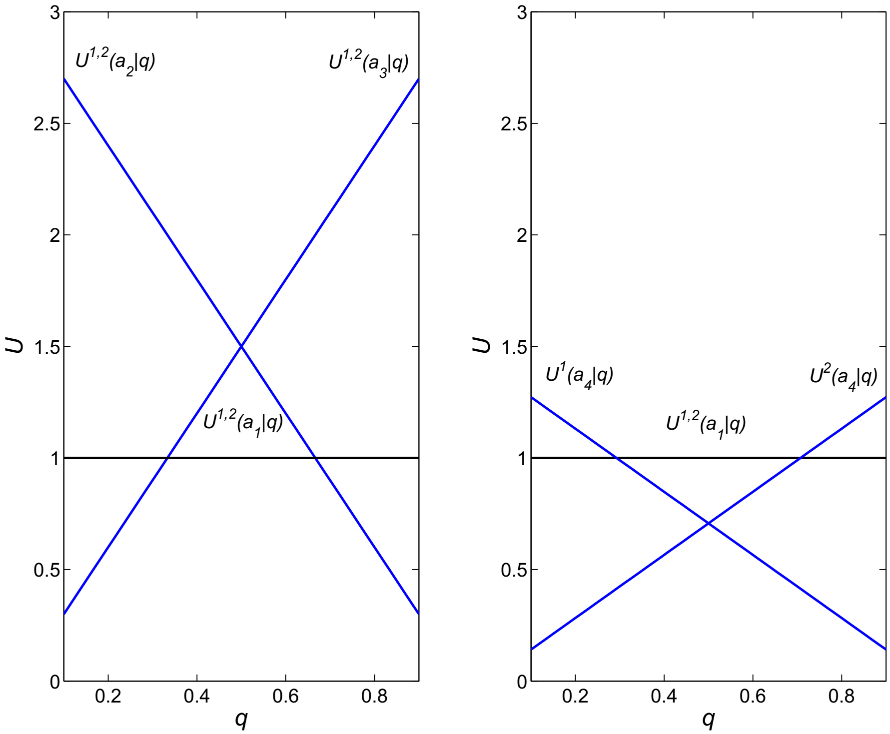

Fig. 2 Expected utilities of agents as a function of q. Left panel: comparing utilities of allocations a_1, a_2 and a_3. Right panel: comparing utilities of allocations a_1 and a_4

whereas the sloped (blue) lines represent the utilities of the two agents under the risky allocation a_4. For low and high q's, one of the agents is better off under a_4 than under a_1, whereas the other is worse off, and for q close to $\frac{1}{2}$ both agents are worse off. Now, the reason one agent is better off for extreme q's is exactly because of speculative redistributions. Regardless of q, allocation a_4 is therefore inferior when using any measure that forces q to be the same for the two agents. What changes with q under allocation a_4 is which one of the two agents reaps the benefits from speculation. The planner can therefore always improve upon a_4.

It follows from the efficiency definitions in the previous section that a_1, a_2, and a_3 are all IK-efficient ($IKE = \{a_1, a_2, a_3\}$), that a_1 and a_4 are belief neutrally inefficient ($WBNE = \{a_2, a_3\}$), that there are no belief neutrally efficient allocations ($BNE = \emptyset$), that $BCEST = \{a_1\}$[13], and that all four allocations are U-efficient

[13] In this example, a_1 is the only maximal allocation regardless of which utility aggregator is used. For a general utility aggregator, $\hat{\mathcal{W}}$, Blume et al. (2018) require that $\hat{\mathcal{W}}(U^1, U^2) \in [\min(U^1, U^2), \max(U^1, U^2)]$, and since $U^1 = U^2$ for allocations a_2 and a_3, for all $q \in \mathcal{Q}_R$, *any* such utility aggregator will rule out these two allocations as being dominated by a_1. For allocation a_4, $U^1(a_4|0.5) = U^2(a_4|0.5) < U^{1,2}(a_1|q)$, for all $q \in \mathcal{Q}_R$, so $\min_{q\in\mathcal{Q}_R} \hat{\mathcal{W}}(U^1(a_4|q), U^2(a_4|q)) < \min_{q\in\mathcal{Q}_R} \hat{\mathcal{W}}(U^1(a_1|q), U^2(a_1|q))$ and thus a_4 is also dominated by a_1. It therefore follows that *any* utility aggregator leads to the unique maximal allocation a_1 in this example.

($UE = \{a_1, a_2, a_3, a_4\}$). Also, it follows that Conditions 4 and 5 are satisfied, whereas Condition 1 is not. These results are robust to allowing for transfers between agents.

The results are consistent with our analysis in the previous section, in that $BNE \subsetneq WBNE \subsetneq IKE \subsetneq UE$, with each inclusion being strict. Specifically, since the utility possibility set is not convex, it is possible for an allocation to be IK-efficient but not weakly belief neutrally efficient (allocation a_1 in this example), since the conditions for Proposition 4 are not satisfied.

One can argue that a_1 should be viewed as efficient in this example. As shown in the left panel of Fig. 2, a_1 is dominated by some other allocation for every $q \in \mathcal{Q}_R$, and is thereby belief neutrally inefficient. However, since the planner with incomplete knowledge about q cannot determine which allocation of a_2 and a_3 to choose over a_1, the allocation is still a reasonable one for the planner who has incomplete knowledge about q.

Note that the measure defined in (16,17) compares the worst case utilities associated with allocations a_2 and a_3 with the constant utility associated with a_1. These allocations are then ruled out even though each dominates a_1 for some reasonable probabilities. As a consequence, only a_1 belongs to *BCEST*. This will be the case even if the a_1 line in the left panel of Fig. 2 is moved downward so that it is barely above the lines associated with a_2 and a_3 at the endpoints. The *BCEST* set is therefore based on a very conservative measure.

If objective randomization is possible, a_1 no longer remains IK-efficient, because a randomization of a_2 and a_3 with equal probability leads to expected utility of 1.5 for both agents in both states, regardless of q. In this case, it therefore follows that $WBNE = IKE = \{a_2, a_3\}$, in line with Propositions 4 and 5, since the utility possibility set is convex, i.e., Condition 1 is satisfied when objective randomization is possible.

4.3 A production economy with non-convex reasonable belief set

The previous example explored an economy in which the utility possibility set was not convex. Here, we explore the consequences of having a non-convex reasonable belief set.

Consider a one-date economy with two agents and three mutually exclusive production technologies that depend on the outcome of three tosses of a coin. If the outcome of the tosses is three tails, production technology a_2 delivers one unit of utility to each agent, otherwise 0. If the outcome is three heads, production technology a_3 delivers one unit of utility to each agent, otherwise 0. If the outcome is neither 3 heads, nor 3 tails, production technology a_4 delivers one unit of utility to each agent, otherwise 0. The two agents agree that the three tosses are independent and identically distributed, but not on the probability, p, for heads in each toss, believing it is p^1 and p^2, respectively, where we assume that $p^1 < p^2$. The economy allows for objective randomization, and the example is also robust to allowing for transfers.

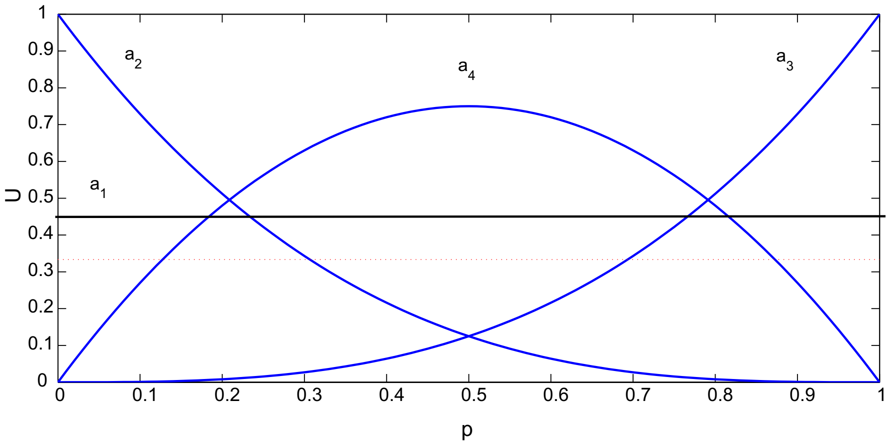

Fig. 3 Utilities associated with production technologies as a function of probability for heads, p. The curves show utilities for the risky technologies, $a_2 - a_4$, whereas the straight line represents utilities for the risk free technology, a_1. The dotted (red) straight line shows utilities from a randomization of the three risky technologies (colour figure online)

The above probability structure may be viewed as a stylized model for a process where the outcome depends on a sum of i.i.d. random variables. For instance, the three technologies could represent different types of seeds that a farmer can choose between. Let the states represent rain or sun. Seeds of type 2 (corresponding to a_2 above) require three days of sun, seeds of type 3 (corresponding to a_3) require three days of rain and seeds of type 4 require at least one day of rain and one day of sun.[14]

Since agents agree on the i.i.d. nature of the coin tosses but disagree on p, it is natural for the planner to include all probability vectors for the states

$$\{\text{ All heads , All tails , Both heads and tails }\}$$

on the form

$$(q_1, q_2, q_3) = (p^3, (1-p)^3, 1 - p^3 - (1-p)^3),$$

for $p \in [p^1, p^2]$. Note that this corresponds to a view that either agent can be correct in his or her belief about the probability, and that probabilities in-between the agents' individual beliefs about individual coin toss are also reasonable. The corresponding set $\mathcal{Q}_R$ is obviously not the convex hull of the two agents' beliefs, which corresponds to

[14] Similar situations arise naturally in biology and in physics. In epidemiology, the reproductive ratio represents the number of individuals infected by a single individual, which determines whether a virus spreads among the population. In nuclear physics, the radioactive decay rate of an atom's nucleus determines the success of fission.

$$(q_1, q_2, q_3) \in CH\Big(\big\{((p^1)^3, (1-p^1)^3, 1-(p^1)^3-(1-p^1)^3),\\ ((p^2)^3, (1-p^2)^3, 1-(p^2)^3-(1-p^2)^3)\big\}\Big).$$

In fact, it is easy to check that $\mathcal{Q}_R$ is not even convex.

The convex hull corresponds to the planner taking the view that there is some probability that agent 1 is correct about p, and that otherwise agent 2 is correct, but neglecting the possibility that some p in-between their beliefs for heads may actually be correct. The convex hull represents such mixtures of the two probability vectors.[15]

The utilities associated with the three technologies are shown in Fig. 3 (blue lines), as a function of p. In addition to the three risky technologies there is a risk-free technology a_1 that delivers a utility of 0.45 to all agents (the black straight line). We also show the utility of a randomization with equal probabilities for the three risky production technologies (as represented by the dotted red straight line).

One may argue that a_1 should be viewed as efficient. Indeed, it is easy to check that a_1 is IK-efficient, since it is above any other technology (including all randomizations) for some p. However, a_1 is belief neutral inefficient, since it is below some other technology for each p. This is consistent with Proposition 5, since Condition 4 is not satisfied in this economy, i.e., the reasonable belief set is not convex.

4.4 A production economy in which *IKE* $\subsetneq$ *WBNE*

Consider only allocations a_2 and a_3 in the example in Sect. 4.2, with reasonable belief set, $\mathcal{Q}_R = \{(q_1, 1-q_1) : q_1 \in [0.2, 0.5]\}$. Both allocations are efficient in the homogeneous beliefs economy with $q_1 = 0.5$, and both allocations therefore belong to *WBNE*. However, a_3 is IK-inefficient, since it is dominated by a_2 for all q in $\mathcal{Q}_R$ except when $q = (0.5, 0.5)$, and moreover does not dominate a_2 for any q in $\mathcal{Q}_R$. One could argue that excluding a_3 from the set of efficient allocations in this example is indeed appropriate, but since $WBNE = \cup_{q\in\mathcal{Q}} E_q$, there is no possibility to exclude an allocation that belongs E_q for some $q \in \mathcal{Q}_R$ because it is inferior for some (all) other $q' \in \mathcal{Q}_R$.

4.5 A production economy in which *IKE* $\subsetneq$ *BCEST*

Consider a modification of the example in Sect. 4.2, in which there are two allocations in the production set: a_3 and another allocation a_5 that pays 0.3^2 in both states. Let the reasonable belief set as before be $\mathcal{Q}_R = \{(q_1, 1-q_1) : q_1 \in [0.1, 0.9]\}$. In this case, $BCEST = \{a_3, a_5\}$ as both allocations yield utility of 0.3 for the worst case scenario of $q = 0.1$. However, $IKE = \{a_3\}$ as allocation a_5 is dominated by a_3 for any $0.1 < q \leq 0.9$.

[15] The problem described here of choosing an appropriate set, $\mathcal{Q}_R$, is more generally related to that of defining appropriate sets in dynamic multiple priors models, as, e.g., analyzed in Epstein and Schneider (2003).

4.6 A production economy with an allocation that is *BCEST* but not *WBNE* nor *IKE*

Again, we modify the example in Sect. 4.2, this time by adding an additional allocation a_6 in the aggregate production set. Allocation a_6 pays 2 to each agent in state 2 and $\left(\frac{1}{0.9}\left(1-0.1\sqrt{(2)}\right)\right)^2$ in state 1. Again, the reasonable belief set is $\mathcal{Q}_R = \{(q_1, 1-q_1) : q_1 \in [0.1, 0.9]\}$. In this modified example, the expected utility when $q = 0.9$ is 1, which is the same as a_1. However, for every other q the utility of a_6 is strictly higher than that of a_1. Hence, allocation a_1 is no longer in *IKE*, as it is dominated by the new allocation a_6 for some reasonable beliefs and never dominates that allocation. Moreover, neither a_1 nor a_6 belong to *WBNE*, since for each reasonable q, either a_2 or a_3 or both dominate each of these allocations. Figure 7 in the Appendix shows the expected utilities in this modified example.

Allocation a_1 belongs to *BCEST*, since the expected utility for both agents is the same under a_6 as under a_1 when $q = 0.9$. Both a_2 and a_3 are inefficient according to *BCEST*, just as in the example in Sect. 4.2. Hence, we have that $IKE = \left\{a_2, a_3, a_6\right\}$, $WBNE = \left\{a_2, a_3\right\}$, $BCEST = \left\{a_1, a_6\right\}$ for this modified example, so *BCEST* is different from both *IKE* and *BNE* and is, moreover, the only criteria that considers a_1 as efficient.

4.7 A limitation of multiple-beliefs based efficiency measures

We conclude with an example that suggests that multiple beliefs-based efficiency measures—as *BNE*, *WBNE*, *IKE*, and *UE* all are—may not work well if one moves even further away from the standard economic environments.

There are two agents, one of whom susceptible to a flu-virus whereas the other is resistant. There is one very potent vaccine dose, that can be distributed to one of the two agents, or be discarded. The agents disagree about who is most likely to be susceptible. Agent $i \in \{1, 2\}$ thinks the probability that (s)he is the susceptible one is $z^i < \frac{1}{2}$. The probability for an unvaccinated susceptible agent to catch the flu is 0.9, and for an unvaccinated resistant agent it is 0.1, facts that both agents agree on. A vaccinated agent cannot catch the flu.

Both agents would like to avoid catching the flu, but also has a high ethical standard, so neither agent would ever consider taking the vaccine unless they are at least as likely to be susceptible as the other agent, i.e., if $z^i \geq \frac{1}{2}$. Objective randomization is not possible.

Agent i's utility when avoiding catching the flu, if unvaccinated or if vaccinated when $z^i \geq \frac{1}{2}$, is 1. The agent's utility associated with catching the flu is 0. Finally, an agent associates utility -1 with being vaccinated when $z_i < 0.5$, because of the ethical issue this creates. Both agents believe $z^i < \frac{1}{2}$, $i = \{1, 2\}$, so they disagree, since under agreement $z^1 + z^2 = 1$.

One verifies that under agreement, it is efficient for the planner to vaccinate the agent with the highest z^i (or either one of the agents in the knife-edge case $z^1 = z^2 = \frac{1}{2}$). However, a vaccinated agent who believes $z^i < \frac{1}{2}$ will actually be worse off than if unvaccinated, regardless of whether (s)he comes down with the flu or not. One may therefore argue that the efficient outcome in this case is to dispose the vaccine dose, since neither agent wants to take it.[16] The multiple-beliefs based measure is unable to capture this concern, since it will always base the welfare function on a probability measure for which $z^1 + z^2 = 1$.

The difference between this example and the previous ones is that beliefs here directly affect utilities. In contrast, in the previous examples, utility was derived from consumption and, crucially, ex post high consumption realizations were valued highly by agents regardless of their prior beliefs about the likelihood of such realizations. It is unclear whether the beliefs-based efficiency measures analyzed in this paper can be applied within these type of economic environments.

A. Proofs

Proof of Proposition 1: *(i)* It immediately follows from (12,13) that $E_q = BNE = WBNE$ in the agreement economy. From (8) it follows that under agreement, an IK-efficient allocation, a satisfies

$$\exists \lambda \in S^N, \forall b \in \mathcal{A}, a \geq_q^\lambda b, \tag{22}$$

which is a stronger condition than (10), since λ is allowed to depend on b in (10) but not in (22), so $IKE \subset BNE$. Also, we note that (8) implies that

$$IKE = \cup_{\lambda \in S^N} M_{q,\lambda}$$

in the agreement economy. Moreover, when $\mathcal{Q}_R = \{q\}$, $\mathcal{W}(a, \lambda) = q'\mathcal{V}(a)\lambda$, so $\arg\max_{a\in\mathcal{A}} \mathcal{W}(a, \lambda) = M_{q,\lambda}$, and thus $BCEST = \cup_{\lambda \in S^N} M_{q,\lambda} = IKE$.

Finally, (9) is obviously equivalent to (14) when $\mathcal{Q}_R = \{q\}$, so $WBNE = UE$.

(ii) From (i), it is sufficient to show that $E_q \subset IKE$ when the utility possibility set is convex and the economy allows for transfers. Again, we note that $IKE = \cup_{\lambda \in S^N} M_{q,\lambda}$ in the agreement economy. Define $Z_q = \{q'\mathcal{V}(a) : a \in \mathcal{A}\}$, which is compact in $\mathbb{R}^N$, and convex when $\mathcal{U}$ is convex. Any Pareto efficient allocation given q, $a \in E_q$, must be associated with a $z_a = q'\mathcal{V}(a)$ on the boundary of Z_q. By the supporting hyperplane theorem, $z_a'\lambda = K$, and $z'\lambda \leq K$ for all $z \in Z_q$, for some $\lambda \in R^N$ and $K \in \mathbb{R}$. Moreover, since a is efficient, $\lambda^n \geq 0$, $n = 1, \dots, N$, with at least one coefficient being strictly positive. This in turn implies that $\lambda \in \bar{S}^N$ can be chosen, and it then follows that

$$E_q \subset \cup_{\lambda \in \bar{S}^N} M_{q,\lambda}.$$

[16] We use the analogy of a virus and vaccine in this example, but there are many other situations where similar considerations may affect an agent's utility, for example, as a charity recipient or selected winner of a prize. In our example, if it would be possible to educate the agents so that they agreed on who was most susceptible, or to convince them to side-step their ethical concerns in the interest of not wasting a viable vaccine dose, this would lead to a superior outcome. We do not consider such alternative possibilities here, in line with the rest of the literature.

To infer that

$$E_q \subset \cup_{\lambda \in S^N} M_{q,\lambda} = IKE,$$

we show that given an allocation $a \in M_{q,\lambda} \cap E_q$ for some $\lambda \in \bar{S}^N$ on the boundary of the unit simplex, there exists a $\hat{\lambda} \in S^N$, such that $a \in M_{q,\hat{\lambda}}$. This then implies

$$E_q \cap (\cup_{\lambda \in \bar{S}^N} M_{q,\lambda}) \subset \cup_{\lambda \in S^N} M_{q,\lambda} = IKE,$$

and $E_q \subset IKE$ then immediately follows.

Assume $a \in E_q \cap M_{q,\lambda}$, so that $z_a' \lambda = K$, for some $\lambda \in \bar{S}^N$, where $\lambda^n = 0$ for agents $n \in B$ in the nonempty set $B \subset \{1, 2, \dots, N\}, 1 \le |B| \le N - 1$. Define $\ell = \min_{k \notin B} \lambda^k$, so that $0 < \ell \le 1$. Agents in B are then allocated zero units of the good $a_{t,m,n} = 0$ when $n \in B$, since they have zero Pareto weight. Also, define

$$\alpha = \max_{t,m,n} u^n_{t,m}(0),$$
$$\beta = \inf_{t,m,n,a \in \mathcal{A}} u^n_{t,m}(a_{t,m,n}),$$
$$\gamma = \frac{\beta}{\alpha}.$$

Because of the compactness of $\mathcal{A}$ and the continuous differentiability of the utility functions, it follows that $0 < \gamma \le 1$. Define $\hat{\lambda} \in S^N$, such that $\hat{\lambda}^n = \frac{\ell\gamma}{N}$, for $n \in B$, and $\hat{\lambda}^k = \left(1 - \frac{|B|\ell\gamma}{N}\right)\lambda^k$ for $k \notin B$.

For $n \in B$, $k \notin B$,

$$\begin{aligned}
\frac{\hat{\lambda}^n}{\hat{\lambda}^k} &= \frac{\frac{\ell\gamma}{N}}{\left(1 - \frac{|B|\ell\gamma}{N}\right)\lambda^k} \\
&= \frac{\ell}{\lambda^k} \frac{1}{N\left(1 - \frac{|B|\ell\gamma}{N}\right)}\gamma \\
&\le \frac{1}{N - |B|\ell\gamma}\gamma \\
&\le \gamma
\end{aligned}$$

Thus,

$$\begin{aligned}
\hat{\lambda}^n &\le \gamma \hat{\lambda}^k \\
&\le \frac{(u^n_{t,m})'(0)}{(u^k_{t,m})'(a_{t,m,n})}\hat{\lambda}^k,
\end{aligned}$$

for all t and m, and so

$$\hat{\lambda}^n (u^n_{t,m})'(0) < \hat{\lambda}^k (u^n_{t,m})'(a_{t,m,k}), \tag{23}$$

for all $n \in B$, and $k \notin B$. As a consequence $a \in M_{q,\hat{\lambda}}$. Specifically, any positive allocation to an agent in B, at any time, in any state, could by (23) be improved upon by allocating the good away from the agent to agent k. Such a transfer would be feasible since the economy allows for transfers.

Finally, if $a \in M_{q,\lambda}$ then it must be that $a \in M_{q,\hat{\lambda}}$, since $z'\hat{\lambda} = \left(1 - \frac{|B|\ell\gamma}{N}\right)z'\lambda$ for any $z \in \mathbb{R}^N$ such that $z_n = 0$ for $n \in B$, and thus the same set of $z_a \in Z_q$ among such zs maximize $z_a'\lambda$ and $z_a\hat{\lambda}$. This completes the proof. □

Proof of Proposition 2 *(i)* Assume $a \notin UE$, in which case there exists a b such that for all $\lambda \in S^N$ and $q \in \mathcal{Q}_R$,

$$\begin{array}{lll} \sum_n \lambda^n U^n(b|q) & > \sum_n \lambda^n U^n(a|q) & \Rightarrow \\ \sum_n \lambda^n U^n(b|q^k) & > \sum_n \lambda^n U^n(a|q^k),\ k = 1, 2, \dots, N & \Rightarrow \\ U^k(b|q^k) & > U^k(a|q^k),\ k = 1, 2, \dots, N & \Rightarrow \\ \sum_n \hat{\lambda}^n U^k(b|q^k) & > \sum_n \hat{\lambda}^n U^n(a|q^n),\ \forall \hat{\lambda} \in S^N & \end{array}$$

and thus $a \notin E^A_{\mathbf{q}}$, so the result follows. Note that it is crucial that $q^k \in \mathcal{Q}_R$ for $k = 1, 2, \dots, N$ for the argument.

(ii) Assume $a \in E^E_{\mathbf{q}}$, and (without loss of generality) that agents 1 and 2 disagree about the relative likelihood of states 1 and 2, such that $\frac{q_1^1}{q_2^1} > \frac{q_1^2}{q_2^2}$, and that both agents are allocated strictly positive amounts of the good in both these states. Define the aggregate production $X = \mathcal{P}(a)$, and note that the argument in Propositions 1 and part (ii) of this proposition above imply that $a \in M_{\mathbf{q},\lambda|X} = \arg\max_{a \in \mathcal{A}, \mathcal{P}(a) = X} \sum_n \sum_m q_m^n [\mathcal{V}(a)]_{m,n} \lambda^n$, for some $\lambda \in S^N$ (since the singleton $\mathcal{A}_X = \{X\}$ is convex).

Because transfers are possible, the following F.O.C. for the allocation needs to be satisfied:

$$\frac{q_1^1 (u_{t,1}^1)'(a_{t,1,1})}{q_2^1 (u_{t,2}^1)'(a_{t,2,1})} = \frac{q_1^2 (u_{t,1}^2)'(a_{t,1,2})}{q_2^2 (u_{t,2}^2)'(a_{t,2,2})} = \frac{\lambda^1}{\lambda^2}. \tag{24}$$

However, for any $q \in \mathcal{Q}_R$, the same argument implies that for $a \in E_q$ with strictly positive allocations to both agents in both states,

$$\frac{(u_{t,1}^1)'(a_{t,1,1})}{(u_{t,2}^1)'(a_{t,2,1})} = \frac{(u_{t,1}^2)'(a_{t,1,2})}{(u_{t,2}^2)'(a_{t,2,2})} = \frac{\hat{\lambda}^1}{\hat{\lambda}^2}, \tag{25}$$

for some $\hat{\lambda} \in S^N$, so $a \notin E_q$ for any such q, since $\frac{q_1^1}{q_2^1} > \frac{q_1^2}{q_2^2}$. Thus $a \notin WBNE$.

Similarly, for all $\hat{\lambda} \in S^N$, the allocation a which satisfies (24) is improved upon by allocation b, in that $b >_q^{\hat{\lambda}}$, that redistributes between agents 1 and 2, such that (25) is satisfied for $\frac{\hat{\lambda}^1}{\hat{\lambda}^2}$ similar to $\frac{\lambda^1}{\lambda^2}$. If $\frac{\hat{\lambda}^1}{\hat{\lambda}^2}$ is sufficiently large or small, an improvement is achieved by allocating the good completely to one of the agents in one or both of the two states states. It follows from (7) that $a \notin IKE$. □

Proof of Proposition 3 Since the economy allows for transfers and the utility possibility set is convex, it follows from Proposition 1 that $E_q = \cup_{\lambda \in S^N} M_{q,\lambda}$. Moreover, from (12),

$$WBNE = \cup_{q \in \mathcal{Q}_R} E_q.$$

If $|M_{q,\lambda}| = 1$ for all $q \in \mathcal{Q}_R$, and $\lambda \in S^N$, the strict dominance condition then follows immediately for all $a \in M_{q,\lambda} \subset WBNE$ (with the associated λ and q defining the parameters in the dominance condition).

For $a \in M_{q,\lambda}$, assume that $M_{q,\lambda}$ contains another element $b \neq a$, i.e., that $U(a|q,\lambda) = U(b|q,\lambda)$. Since the aggregate production set is convex, $c = \frac{1}{2}a + \frac{1}{2}b \in \mathcal{A}$. Moreover, from (3) and the strict concavity of agents' utilities, it follows that $U(\cdot|q,\lambda)$ is strictly concave over allocations, and thus

$$U(c|q,\lambda) > \frac{1}{2}U(a|q,\lambda) + \frac{1}{2}U(b|q,\lambda) = U(a|q,\lambda),$$

contradicting the assumption that $a \in M_{q,\lambda}$. So, no such $b \neq a \in E_{q,\lambda}$ exists, $|E_{q,\lambda}| = 1$, and the result therefore follows. □

Proof of Proposition 4 We show that $a \notin WBNE \Rightarrow a \notin IKE$. For an arbitrary $\lambda \in S^N$, define the mapping $\mathcal{F}_\lambda : \mathcal{U} \to \mathbb{R}^M$, by $\mathcal{F}_\lambda(V) = V\lambda$, and the set $F_\lambda = \mathcal{F}_\lambda(\mathcal{U})$, which is a convex, compact, subset of $\mathbb{R}^M$, because of Condition 1 and the compactness of the production set. Moreover, define $f_a = \mathcal{F}_\lambda(\mathcal{V}(a))$. Since $a \notin WBNE$, it follows from (9) that

$$\max_{q \in \mathcal{Q}_R} \min_{f \in F_\lambda} q'(f_a - f) = s < 0,$$

where $s < 0$ follows from the fact that the optimum is realized for some f^*, q^* (because both $\mathcal{Q}_R$ and F_λ are compact).

It follows from Sion's minmax theorem that

$$\min_{f \in F_\lambda} \max_{q \in \mathcal{Q}_R} q'(f_a - f) = s,$$

where the same, f^*, q^* can be chosen for the maxmin and minmax problems, and thus that for all $q \in \mathcal{Q}_R$,

$$q'(f_a - f^*) \leq s < 0,$$

i.e.,

$$q'\mathcal{V}(b_*)\lambda > q'\mathcal{V}(a)\lambda,$$

for all $q \in \mathcal{Q}_R$, for the allocation b_*, such that $b_* = \mathcal{V}(f_*)$. This is, in turn, equivalent to $b_* \geq_q^\lambda$ for all $q \in \mathcal{Q}_R$. Since $\lambda \in S^N$ was arbitrary, it then follows that (7) is satisfied, and thus $a \notin IKE$. □

Proof of Proposition 5 From (12) and Proposition 1, it follows that

$$WBNE = \cup_{q \in \mathcal{Q}_R} \cup_{\lambda \in S^N} M_{q,\lambda}.$$

Consider $a \in WBNE$, and an associated $q \in \mathcal{Q}_R$ and $\lambda \in S^N$. Because of the strict dominance condition, it follows that $\forall b \neq a$, $a >_q^\lambda b$. But this immediately implies (8) for λ and q, where q does not even depend on b in this case, i.e., when $a \in M_{q,\lambda}$ is the only element in that set:

$$\forall b \in \mathcal{A} : a >_q^\lambda b,$$

which immediately implies (8). □

Proof of Proposition 6 Necessity: From (13),

$$BNE = \cap_{q \in \mathcal{Q}_R} E_q.$$

If $a \in E_q$ for all $q \in \mathcal{Q}_R$, then $\mathcal{P}(a) \in \mathcal{P}(E_q)$ for all $q \in \mathcal{Q}_R$. Thus, if $\cap_{q \in \mathcal{Q}_R} \mathcal{P}(E_q) = \emptyset$, then $BNE = \emptyset$.

Sufficiency: From (13) and Proposition 1, it follows that

$$BNE = \cap_{q \in \mathcal{Q}_R} \cup_{\lambda \in S^N} M_{q,\lambda},$$

where $M_{q,\lambda} = \arg\max_{a \in \mathcal{A}} q' \mathcal{V}(a) \lambda$.

We decompose the problem into aggregate production and allocation parts:

$$\max_{a \in \mathcal{A}} q' \mathcal{V}(a) \lambda \max_{X \in \mathcal{A}_X} \max_{a \in \mathcal{A} | \mathcal{P}(a) = X} \sum_n \lambda^n \sum_t \sum_m u^n_{t,m}(a_{t,m,n}) q_m, \tag{26}$$

and focus on the optimization conditional on X. Specifically, we define:

$$M^X_{q,\lambda} = \underset{a \in \mathcal{A} | \mathcal{P}(a) = X}{\arg\max} \sum_n \lambda^n \sum_t \sum_m u^n_{t,m}(a_{t,m,n}) q_m.$$

The optimality condition for the conditional problem, $M^X_{q,\lambda}$, w.r.t. allocations across agents and states are:

$$\lambda^n (u^n_{m,t})'(a_{t,m,n}) = \rho^{t,m} \tag{27}$$

across all t, and m, for all agents for which $a_{t,m,n} > 0$, and

$$\lambda^n (u^n_{t,m})'(0) \leq \rho^{t,m} \tag{28}$$

for all agents, states, and times such that $a_{t,m,n} = 0$, where $\rho^{t,m} > 0$ are the Lagrange multipliers that ensure that aggregate consumption is equal to aggregate production in each state, at each time. Here, we require the Lagrange multipliers to be adapted, when viewed as a process $\rho : \mathcal{T} \times \Omega \to \mathbb{R}_+$. These conditions do not depend on q, which immediately implies □

Lemma 1 *In an economy that allows for transfers, $M^X_{q,\lambda}$ does not depend on $q \in S^N$: $M^X_{q,\lambda} = M^X_\lambda$.*

So, if $X \in \mathcal{P}(E_q)$ for all $q \in \mathcal{Q}_R$, then for $\lambda \in S^N$, $M^X_{q,\lambda} \subset M_{q,\lambda}$ for all $q \in \mathcal{Q}_R$, and

$$BNE = \cap_{q \in \mathcal{Q}_R} \cup_{\lambda \in S^N} M_{q,\lambda} \subset \cap_{q \in \mathcal{Q}_R} M^X_{q,\lambda} = M^X_\lambda \neq \emptyset.$$

Proof of Proposition 7: *(i)* The result follows immediately from Lemma 1, and the fact that $E_q = \cup_\lambda M^X_{q,\lambda} = \cup_\lambda M^X_\lambda$ (since $\{X\} = \mathcal{A}_X$ in the exchange economy), which does not depend on q.

(ii) From (i) and (12–13), it follows that $BNE = WBNE$.

Now, $M_{q,\lambda} = M^X_\lambda$, i.e., $M_{q,\lambda}$ does not depend on q. So, since $a \in WBNE \Rightarrow a \in M^X_\lambda$ for some $\lambda \in S^N$, and $a \in M^X_\lambda \Rightarrow \forall b \in \mathcal{A}, \forall q \in \mathcal{Q}_R : a \geq^\lambda_q b$, this implies that condition (8) is satisfied, i.e., $a \in IKE$. Thus, $WBNE \subset IKE$.

For $a \notin WBNE$, it follows that $a \notin M^X_\lambda$ for all $\lambda \in S^N$, i.e., $\forall \lambda \in S^N, \exists b \in \mathcal{A}, \forall q \in \mathcal{Q}_R : b >^\lambda_q a$, which by (7) implies $a \notin IKE$. So, $IKE \subset WBNE$, and thus $IKE = WBNE$.

Finally, we show that $BCEST = WBNE$. The result follows from the following lemma: □

Lemma 2 *Consider a continuous function* $F : \mathcal{A} \times \mathcal{Q}_R \to \mathbb{R}$, *and define the sets* $Y_q = \arg\max_{a \in \mathcal{A}} F(a, q), q \in \mathcal{Q}$. *Also, define the set* $L = \arg\max_{a \in \mathcal{A}} \min_{q \in \mathcal{Q}_R} F(a, q)$. *Then, if* $Y_q = Y_{q'} \overset{\text{def}}{=} Y$ *for all* $q, q' \in \mathcal{Q}_R$, *it follows that* $L = Y$.

Proof of Lemma 2 $L \subset Y$: Assume $s \notin Y$, then $F(y, q) > F(s, q)$ for all q, for some $y \in Y$ (Y is nonempty because of compactness of $\mathcal{A}$ and continuity of F). Define $q^* = \arg\min_{q \in \mathcal{Q}} F(y, q)$. Then, $F(y, q^*) = \min_{q \in \mathcal{Q}_R} F(y, q) > F(s, q^*) \geq \min_{q \in \mathcal{Q}_R} F(s, q)$, so $s \notin L$, and thus $L \subset Y$.

$Y \subset L$: Note that L is nonempty (because of continuity and compactness), so there must exist an $y \in Y \cap L$. Define $z = \min_{q \in \mathcal{Q}_R} F(y, q)$. Now consider any $y' \in Y$. Since $Y_q = Y$ for all $q \in \mathcal{Q}_R$, it follows that $F(y, q) = F(y', q)$, for all q, and therefore $z = \min_{q \in \mathcal{Q}_R} F(y', q)$, which in turn implies that $y' \in L$. So $Y = L$. □

Now, Lemma 2 applied to the function $F_\lambda(a, q) = q'\mathcal{V}(a)\lambda$, and $L_\lambda = \arg\max_{a \in \mathcal{A}} \min_{q \in \mathcal{Q}_R} F_\lambda(a, q)$

$$BCEST = \cup_{\lambda \in S^N} L_\lambda = \cup_{\lambda \in S^N} M^X_\lambda = WBNE,$$

where the middle equality follows from the fact that $M^X_{q,\lambda} = M^X_\lambda$ independently of q in the exchange economy, as implied by Lemma 1.

Thus, altogether, $BCEST = IKE = WBNE = BNE$.

B. Further discussion and comparison of results and proofs

The following list of definitions and conditions are introduced in the paper:

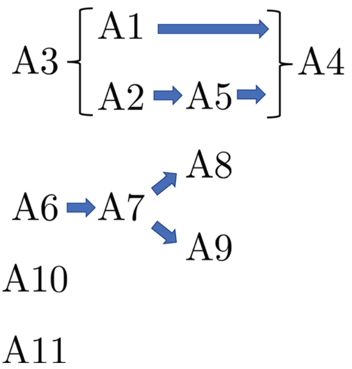

- Proposition 1
  (i) A6
  (ii) A6, A2, A4
- Proposition 2, ¬A6
  (i) A9
  (ii) A11
- Proposition 3, A1,A9⇒A10
- Proposition 4, A1, A4, A8
- Proposition 5, A1, A4, A10
- Proposition 6, A1, A5
- Proposition 7, A3

Fig. 4 Left: relations between definitions and conditions, A1–A11; A3 is equivalent to A1∧A2, A2 implies A4, A1∧A5 implies A4, A6 implies A7, A7 implies A8 and A9. Right: list of conditions needed in propositions

A1. Economy allows for transfers, Definition 1.
A2. Aggregate production set is singleton, $\mathcal{A}_X = \{X\}$.
A3. Exchange economy, i.e., satisfies A1 and A2, Definition 2.
A4. Utility possibility set is convex, Condition 1.
A5. Aggregate production set is convex, Condition 2.
A6. Agents agree, $q^n = q, q = 1, 2, \ldots, N$.
A7. Set of reasonable beliefs is convex hull of individual beliefs, $\mathcal{Q}_R = CH(\{q^n\}_n)$.
A8. Set of reasonable beliefs is convex, Condition 4.
A9. Individual beliefs belong to set of reasonable beliefs, Condition 3.
A10. Strict dominance condition is satisfied, Condition 5.
A11. Agents get positive allocation in states they disagree about, condition in Proposition 2(ii).

The Venn diagram in Fig. 1 holds in general economies, without further assumptions. Figure 4 shows additional assumptions made in Propositions 1–7, as well as the relations between the different concepts in the list above.

Below, we provide a list of results and examples in the paper that show how the efficiency concepts are related in different economic environments under disagreement:

1. In general: $BNE \subset WBNE$, Fig. 1.
2. In general: $WBNE \cup IKE \cup BCEST \subset UE$, Fig. 1.
3. Condition for: $E^A_{\mathbf{q}} \subset UE$, Proposition 2(i).
4. Condition for: Allocation $a \in E^A_{\mathbf{q}}, a \notin (IKE \cup WBNE)$, Proposition 2(ii).
5. Example: $E^A_{\mathbf{q}} \cap UE = \emptyset$, page 13.
6. Condition for: $IKE \subset WBNE$, Proposition 4.
7. Condition for: $WBNE \subset IKE$, Proposition 5.

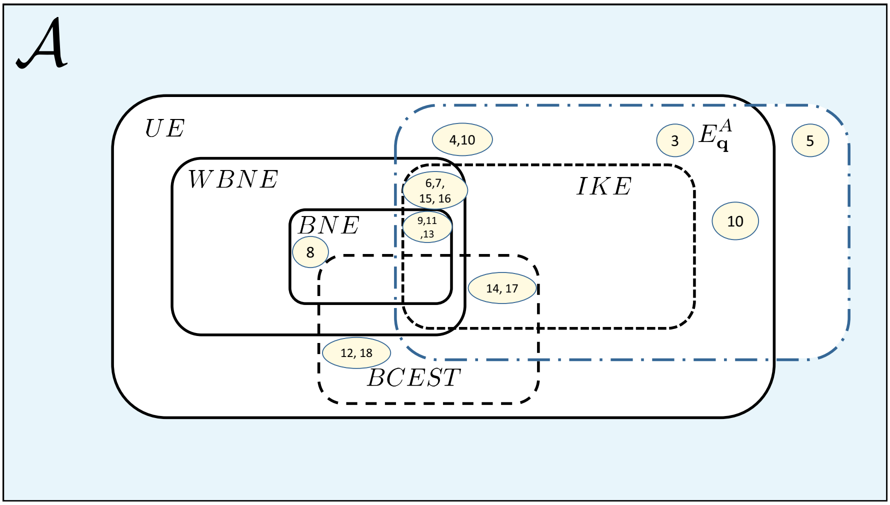

Fig. 5 Conceptual figure, showing how examples and results relate to the Venn diagram in the main paper

8. Condition for: $BNE = \emptyset$, Proposition 6.
9. Exchange economy: $IKE = BCEST = WBNE = BNE$, Proposition 7(ii).
10. Exchange economy: $E_{\mathbf{q}}^{A} \cap IKE = \emptyset$, Proposition 7(iii).
11. Exchange economy example: $WBNE = BNE = IKE = BCEST$, page 16.
12. Example: $BCEST \subsetneq (WBNE \cap IKE)$, production economy, page 18.
13. Example: Standard production economy with $WBNE = IKE$ and $BNE = \emptyset$, Sect. 4.1.
14. Example: Production economy with non-convex utility possibility set: $BNE \subsetneq WBNE \subsetneq IKE \subsetneq UE$, $BCEST \subsetneq IKE$, and $BCEST \cap WBNE = \emptyset$, Sect. 4.2.
15. Example: Production economy with non-convex reasonable belief set : $WBNE \subsetneq IKE$ (robust to randomization), Sect. 4.3.
16. Example: $IKE \subsetneq WBNE$, Sect. 4.4.
17. Example: $IKE \subsetneq BCEST$, Sect. 4.5.
18. Example: Allocation, $a \in BCEST$, $a \notin WBNE$, $a \notin IKE$, Sect. 4.6.

In Fig. 5, we provide a conceptual figure that shows how the different examples and results, identified by the list above, relate to the Venn diagram in the main paper.

The list below provides an informal discussion about the proofs, with a focus on the intuition behind the results:

- Proposition 1: The result follows from the important property of the Pareto efficient allocations in the agreement economy, E_q, that they are characterized by the union of all solutions to the Planner's problem over S^N: $E_q = \cup_{\lambda \in S^N} M_{q,\lambda}$. This allows one to "build" the different efficiency sets from varying beliefs and Pareto weights. Some work goes into showing that the union of Pareto weights only needs to be taken over the open unit simplex, S^N, rather than over its closure, $\bar{S}^N$. Propositions 1(i) and (ii) then follow directly from the definitions of the different efficiency measures, and the fact that when there is only one reasonable belief (because of agreement), $\mathcal{Q}_R = \{q\}$, the measures "collapse" in this dimension and are therefore significantly simplified. Technically, the standard set inclusion result for from predicate calculus for Boolean functions $F : A \times B \to \{F, T\}$ (where A and B are nonempty) are used

$$\forall x, \forall y : F(x,y) = T \Rightarrow \exists x \forall y : F(x,y) = T \Rightarrow$$
$$\forall y \exists x : F(x,y) = T \Rightarrow \exists x \exists y : F(x,y) = T.$$

- Proposition 2: For (i), the result follows because the definition of UE requires an improvement for all $q \in \mathcal{Q}_R$ *and* $\lambda \in S^N$, for an allocation to be considered inefficient. This makes allocations in $E^A_{\mathbf{q}}$ efficient, since they are efficient given individual agent beliefs, which per assumption belong to the set of reasonable beliefs. Condition (ii) highlight the fact that allocations in $E^A_{\mathbf{q}}$ are speculative in the disagreement economy, and therefore deviate from allocations in agreement economies. For this to hold, disagreeing agents must actually be allocated the good, an assumption that is therefore made. Comparing part (i) with part (ii) of the proposition, the power of *WBNE* and *IKE* compared with *UE* is that since they only require "for all" in one of the dimensions q and λ, they are able to rule out speculative allocations as inefficient. This is a main reason why they were developed.
- Proposition 3: The result follows from standard uniquness of minimum to this strictly concave maximization problem.
- Proposition 4: The result follows from the standard duality theory of solutions to max-min and min-max problems, as shown by Sion's theorem (a generalization of the classical minmax theorem). The former problem characterizes *WBNE* solutions and the latter *IKE* solutions, per definition. Note that convexity of the reasonable belief set and the utility possibility set are needed for Sion's theorem to apply.
- Proposition 5: The result is almost immediate, because of the simple characterization of elements in *WBNE* as being Pareto efficient in an agreement economy for a reasonable q. Because of strict dominance, such allocations are unique conditioned on q and λ, and therefore immediately satisfies the condition for *IKE*.
- Proposition 6: The result basically formalizes the intuition that *BNE* is too restrictive a concept for the set to be nonempty, whenever different beliefs lead to different optimal aggregate production levels. Specifically, in economies with standard production sets that allow for transfers, given a level of aggregate production, *BNE* and *WBNE* agree on efficient allocations *between* agents. However, when different reasonable beliefs, q^a and q^b, lead to different optimal aggregate production levels, there is no way to satisfy the requirements of *BNE*. Indeed any efficient allocation conditioned on q^a can be improved if q^b is correct, by changing aggregate production, and vice versa. *BNE* is therefore empty.
- Proposition 7: Like in Proposition 1, where the efficiency concepts "collapsed" in a production economy under agreement, they "collapse" here too, in an exchange economy under disagreement. Specifically, since all concepts are based on efficiency conditioned on a reasonable q that agents hypotheti-

cally agree upon, and under such hypothetical agreement allocations actually do not depend on the actual q in the exchange economy, the result follows.

C. Supporting material for examples in Sect. 4

Section 4.1

We consider a version of the economy in Heyerdahl-Larsen and Walden (2022). Specifically, the production economy has two dates, $t = 1, 2$, and it allows for transfers. There are M possible states, and $N > 1$ agents, who disagree. The state is revealed at $t = 2$, so we require that $a_{m,n,1} = a_{m',n,1}$ for $1 \le m, m' \le M$, for all $a \in \mathcal{A}$. Moreover, we assume that agents have strictly concave utility, i.e., that the functions $u^n_{m,t}$ are strictly concave for all n and m, $t \in \{1, 2\}$, and that Condition 4 is satisfied, i.e., that the set of reasonable beliefs is convex.

There is one unit of a divisible and perishable good that can either be consumed at $t = 1$ or invested in a linear production technology. Each unit invested yields a random, strictly positive, amount, $R \in S^M$, at time $t = 2$, at which point it is consumed. This leads to:

Definition 10 An allocation is feasible, $a \in \mathcal{A}$, with aggregate investment, $I \in [0, 1]$, if

(i) $I = 1 - \sum_{n=1}^{N} a_{1,n,1}$,
(ii) $\sum_{n=1}^{N} a_{m,n,2} \le IR_m$, $m = 1, \dots, M$.

It follows that the aggregate production set, $\mathcal{A}_X \subset \mathbb{R}^{2\times 2}$ is convex, and therefore that Condition 1 is satisfied. From Propositions 3–5, IK-efficiency and weak belief neutral efficiency coincide in this economy,

$$IKE = WBNE.$$

This shows that there are interesting production economies in which these efficiency measures coincide.

Next, we study belief neutral efficiency. Consider the special but important case when all agents have separable power utility across states and time: $u^n_{m,t}(c) = \rho^t \frac{c^{1-\gamma}}{1-\gamma}$, for all n, m, and t, with $\gamma \ne 1$, $\rho > 0$. In this case, BNE will typically be empty. Specifically, it is is well known that for any q that is common among all agents, a representative agent formulation of the social planner's problem exists, in which the utility function of the representative agent is the same as for the individual agents, regardless of the planner's Pareto weights, $\lambda \in \mathcal{S}^N$. Optimal investments, I^* therefore depend on q, but not on λ, $I^* = I^*(q)$ in this planner's problem. As long as there are multiple optimal investment level associated with the set of reasonable beliefs, $I^*(q) \ne I^*(q')$, $q, q' \in \mathcal{Q}_R$, it then follows that BNE is empty. It also follows that the conditions of Proposition 2 (i) and (ii) are satisfied, so all Arrow optima are U-efficient, as well as IK-inefficient and belief neutrally inefficient.

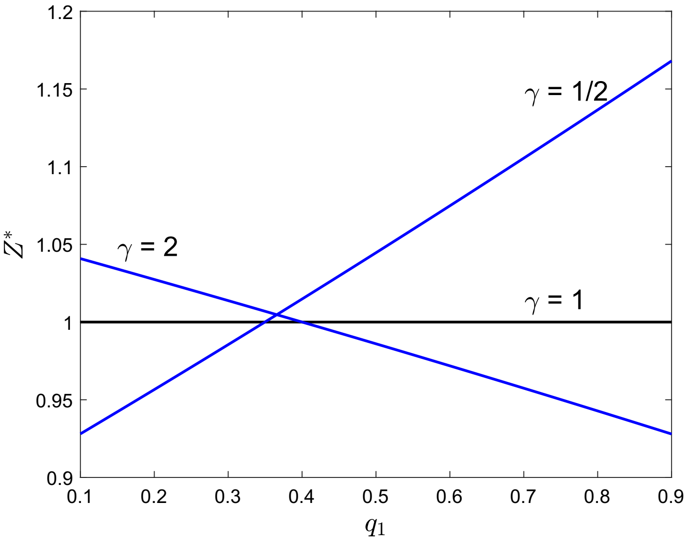

Fig. 6 Optimal investment-to-consumption, as a function of q_1, for different coefficients of risk aversion

To illustrate, let us assume that there are two states with returns $R_1 = 0.9$ and $R_2 = 1.2$. Moreover, assume that there are two agents with beliefs given by $q_1^1 = 0.1$ and $q_1^2 = 0.9$, that the set of reasonable beliefs is the convex hull of the two agents' beliefs, i.e., $\mathcal{Q}_R = \{(q_1, 1-q_1) : 0.1 \leq q_1 \leq 0.9\}$, and that the discount factor is $\rho = 1$.

For an Arrow optimum, $a \in E^A_{(q^1,q^2)}$, with Pareto weights λ^1, λ^2, the first order conditions of optimality then imply that

$$q_m^1 \lambda^1 \left(a_{m,1,2}\right)^{-\gamma} = q_m^2 \lambda^2 \left(a^2_{m,2,2}\right)^{-\gamma}, \; m \in \{1,2\} \qquad \text{and} \qquad \lambda^1 \left(a_{1,1,1}\right)^{-\gamma} = \lambda^2 \left(a_{1,2,1}\right)^{-\gamma},$$

and consequently agent 1 consumes relatively more in state 2 than in state 1, being relatively optimistic about that state. The outcome is therefore speculative. Note that the allocation is U-efficient, since there is no other allocation that is better for both agents for all reasonable beliefs.

Aggregate time 1 consumption is $1 - I$, and therefore the economy's investment-to-consumption ratio is $Z = \frac{I}{1-I}$. For a given q, it is easy to derive the (λ-independent) optimal investment ratio as

$$Z^*(q) = \frac{I^*(q)}{1 - I^*(q)} = E_0\left[\tilde{R}^{1-\gamma}\middle| q\right]^{\frac{1}{\gamma}}. \tag{29}$$

Figure 6 shows the optimal consumption ratio for three different values of risk aversion, $\gamma \in \left\{\frac{1}{2}, 1, 2\right\}$, when varying q_1. When agents have logarithmic utility, corresponding to $\gamma = 1$, the ratio does not vary with q. This corresponds to the well-known

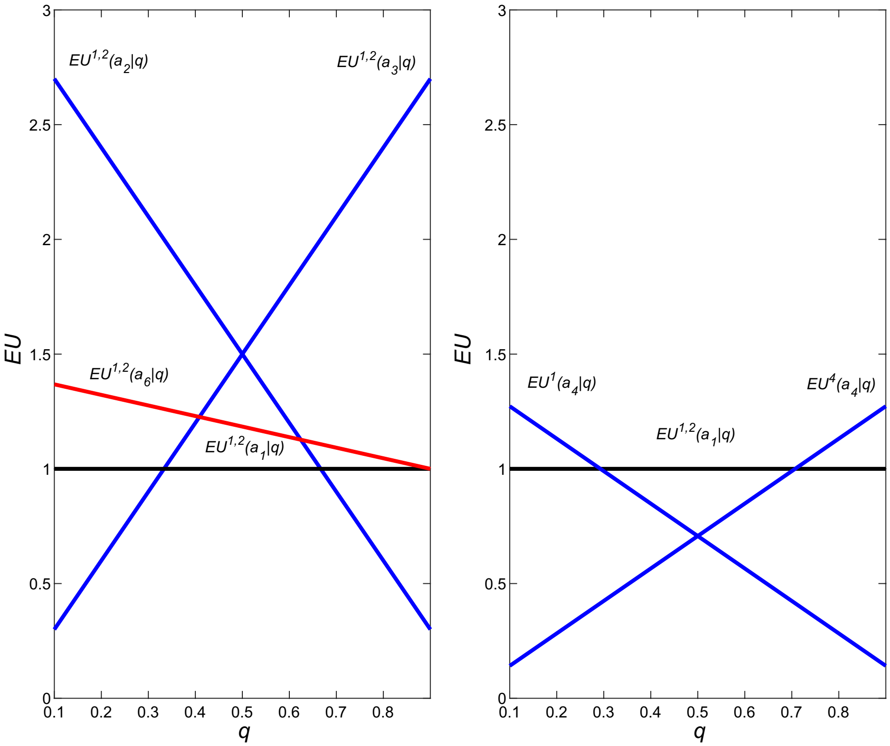

Fig. 7 Modified example in Sect. 4.5. The red line in the left panel represents the expected utility associated with the new allocation a_6, which lies weakly above a_1, and strictly below either a_2 or a_3 for all reasonable beliefs (colour figure online)

property of logarithmic utility that the investment opportunity set does not affect the consumption ratio, see Merton (1969). In this knife edge case, $|I^*| = \frac{1}{2}$, and $BNE = WBNE = IKE$. For risk aversion greater than unity, the optimal consumption ratio increases in q_1 while for risk aversion less than unity the ratio decreases in q_1. Whenever $\gamma \neq 1$, BNE is therefore empty, in line with the discussion above.

Altogether, this work-horse economy example shows that U-efficiency is too broad a concept to rule out speculative allocations, that belief neutral efficiency is such a strong concept that it may rule out all allocations, and that the concept of efficiency under disagreement can be extended to interesting production economies, while capturing both the intuitions provided by weak belief neutral efficiency and IK-efficiency.

Section 4.3

Figure 7 shows expected utilities as function of q for the example in Sect. 4.5. The red line is thae new allocation, a_6, which changes the efficiency sets under the various concepts.

Funding Open access funding provided by Norwegian Business School.

Declarations

Conflict of interest The authors have no relevant financial or non-financial interests to disclose.

References

Bahel E, Sprumont Y (2020) Strategyproof choice of social acts. Am Econ Rev 110:596–627
Blume LE, Cogley T, Easley DA, Sargent TJ, Tsyrennikov V (2018) A case for incomplete markets. J Econ Theory 178:191–221
Brock WA (1979) An integration of stochastic growth theory and the theory of finance, part 1: the growth model. In: Green JR, Scheinkman JA (eds) General equilibrium, growth, and trade. Academic Press, New York, pp 165–192
Brunnermeier MK, Simsek A, Xiong W (2014) A welfare criterion for models with distorted beliefs. Q J Econ 129:1753–97
Cox JC, Ingersoll JE, Ross SA (1985) An intertemporal general equilibrium model of asset prices. Econometrica 53:363–384
Duffie D (2001) Dynamic asset pricing theory. Princeton University Press, Princeton
Dybvig PH (1988) Distributional analysis of portfolio choice. J Bus 61:369–393
Epstein L, Schneider M (2003) Recursive multiple-priors. J Econ Theory 113:1–31
Fleurbaey M (2010) Assessing risky social situations. J Polit Econ 118:649–6808
Gayer G, Gilboa I, Samuelson L, Schmeidler D (2014) Pareto efficiency with different beliefs. J Legal Stud 43:S1541–S1571
Gilboa I, Schmeidler D (1989) Maxmin expected utility with nonunique prior. J Math Econ 18:141–153
Gilboa I, Samet D, Schmeidler D (2004) Utilitarian aggregation of beliefs and tastes. J Polit Econ 112:932–938
Gilboa I, Samuelson L, Schmeidler D (2014) No-betting-pareto dominance. Econometrica 82:1405–1552
Harris R (1978) Ex-post efficiency and resource allocation under uncertainty. Rev Econ Stud 45:427–436
Heyerdahl-Larsen C, Walden J (2022) Distortions and efficiency in production economies with heterogeneous beliefs. Rev Financ Stud 35:1775–1812 (**forthcoming**)
Jouini E, Kallal H (2001) Efficient trading strategies in the presence of market frictions. Rev Financ Stud 14:343–369
Mas-Colell A, Whinston MD, Green JR (1995) Microeconomic theory. Oxford University Press, Oxford
Mehra R, Prescott E (1980) Recursive competitive equilibrium: the case of homogeneous households. Econometrica 48:1365–1379
Merton RC (1969) Lifetime portfolio selection under uncertainty: the continuous time case. Rev Econ Stat 51:247–257
Mongin M (1995) Consistent Bayesian aggregation. J Econ Theory 66:313–351
Mongin M (2016) Spurious unanimity and the Pareto principle. Econ Philos 32:511–532
Starr RM (1973) Optimal production and allocation under uncertainty. Q J Econ 87:81–95
Yaari ME (1981) Rawls, Edgeworth, Shapley, Nash: theories of distributive justice re-examined. J Econ Theory 24:1–39