

Disagreement, Information, and Welfare

Jernej Čopić *

UCLA Economics and Caltech HSS

Abstract

In a stylized strategic situation, two individuals form consistent (self-confirming) assessments as classical statisticians. In equilibrium, where individuals are rational and sophisticated, there are two outcomes: (i) disagreement bears no idiosyncratic risks, minimizes aggregate welfare, individuals cannot recover the truth, and may hold different assessments; (ii) agreement is robust, maximizes welfare, and assessments coincide with the truth. A subjective Pareto criterion compares outcomes based on assessments that players may hold. Whereas agreement is Pareto efficient, disagreement subjectively Pareto-dominates agreement. Under equilibrium assessments, individuals disagree on redistribution. The example relates to “agreeing to disagree” (Aumann 1976), trade and information (Milgrom and Stokey 1982), and a toy macroeconomic example.

*jcopic@econ.ucla.edu. November 5, 2014. A version of this paper has previously been available under the title “Autarky, Exchange, Information, and Welfare.” I am grateful to Andy Atkeson, Tiago Caruso, Aleh Tsivinsky, Romans Pans, Bill Zame, and two anonymous referees, who have commented on that earlier draft.

1 Introduction

It has been argued that in many environments individuals may hold different assessments of the true probability distribution over the uncertain parameters.¹ Barring irrationality, myopia, or inability of economic individuals to process information, such considerations encounter two methodological issues. First, each individual may have a subjective view on what it means to optimize and, in equilibrium, these views are reconciled and tied to the truth. Second, to evaluate different equilibrium outcomes in terms of their welfare, an appropriate welfare notion must also take into account the different assessments that the individuals might possibly hold. In this paper we construct a game to show that it is possible to have two equilibrium outcomes, A and B, where two individuals' welfare evaluations lead them to both prefer A over B, while they also both prefer B over A. The individuals may disagree about the redistributive properties of different equilibrium outcomes. The main contribution here is thus normative.

The following toy macroeconomic example provides an intuitive description of our result. Two individuals can either engage at a lower autarchic level of economic activity, or at a higher level of exchange. The information the individuals can recover from their partial observations of the outcome – each recovers the distribution over her payoffs – is endogenous to the level of economic activity.² Exchange is an equilibrium where there are idiosyncratic risks so that the two individuals can recover the truth perfectly. Exchange results in the highest aggregate welfare. In

¹The term assessment is used to maintain the distinction between the present model, where an assessment is envisioned as a result of some asymptotically consistent estimation procedure, and Bayesian models, where a prior belief is exogenously imposed and need not be consistent with a given economic outcome. Bayesian models with different priors date at least back to Miller 1977 and Harrison and Kreps 1978 in the context of trade. Morris 1995 argues for such models in a game-theoretic setting. Recent such models in macroeconomics are Fostel and Geanakoplos 2008, Geanakoplos 2010, Simsek 2013, and Angeletos and La'o 2013. See Kurz 2011, for a survey of dynamic macroeconomic models with agents' heterogeneous beliefs.

²Uncertainty here is regarding the states of the economy – there is no model uncertainty. For that reason, we use the term “recover” rather than “identify” to not cause any confusion with the standard use of the latter term in Econometrics.

autarky there are no idiosyncratic risk so that the individuals cannot recover the truth.³ Each can support autarky in equilibrium only by overweighing the states that would in exchange be unfavorable to her. Such assessments are consistent with the individuals' observations, and each can justify the other's assessments, and thus the other's behavior. Autarky results in the lowest aggregate welfare and in autarky, the players disagree about the redistributive properties of equilibrium outcomes. Autarky yields a (subjectively) higher welfare than exchange to each individual under any supporting assessment that she might hold. Under the truth, exchange yields a higher welfare than autarky. We describe this example in more detail in Section 6.

We develop our point in the context of an abstract two-player game and then relate it to classical results. In that abstract game, we call the two equilibrium outcomes *agreement* and *disagreement*. In agreement, the two players' assessments must coincide with the truth,⁴ and in disagreement their assessments must differ. These names also suggest a connection with Aumann's 1976 theorem, which is non-coincidental. Agreement has an interpretation, which corresponds tightly to Aumann's theorem: player's assessments are common, their conditional assessments (after receiving a signal) are common knowledge, and must therefore be the same. In disagreement, players' assessments are not common, it is as if their conditional assessments were common knowledge, but these are nonetheless different. Additionally, because of their subjective evaluations of redistributive properties, in disagreement the players cannot jointly agree to switch to agreement. These statements comple-

³The intuition that the information the individuals can recover is endogenous to the level of economic activity is related to Ordoñez 2013, who studies a dynamic economy with learning. Here, we study a static economy, and instead of Bayesian decision makers who have prior beliefs, the individuals are imagined as classical statisticians, who form consistent assessments over uncertainty from their partial observations of economic outcomes and then act rationally.

⁴We say that agreement is *fully recoverable*, in the spirit of the literature on Rational Expectations Equilibria, see e.g., Radner 1979, 1982. There, an outcome is said to be fully revealing if the individuals can deduce the state of the world from their private information and prices, i.e., if all the information is aggregated. In the sense that all the information is aggregated, agreement is also fully revealing. More generally, if an outcome is fully revealing, then it is fully recoverable, but an outcome may be fully recoverable and not be fully revealing.

ment Aumann's 1976 theorem in an equilibrium setting.

Our example also relates to Milgrom and Stokey 1982 no-trade result. In Section 5 we interpret the two equilibrium outcomes as *no-trade* (where assessments coincide) and *trade* (where assessments differ). We reformulate the payoffs in our game so that both outcomes are now *ex-ante* Pareto-efficient under the truth. According to Milgrom and Stokey 1982, when risk-averse traders begin at a Pareto-efficient allocation and have concordant posteriors, they can never agree on a non-null trade in a fully recoverable equilibrium. The no-trade equilibrium tightly corresponds with this result: the individuals' equilibrium assessments (analogous to priors) are common, after receiving their signals about the state of the economy, their conditional assessments (analogous to posteriors) are common knowledge and coincide. Nevertheless, trade is also a possible equilibrium outcome: players' conditional assessments may be common knowledge but, because their assessments are different, these conditional assessments do not coincide. Even while the initial allocation is Pareto-efficient, in an equilibrium outcome where traders' information is different, there can be trade. More generally, if the initial allocation is Pareto-efficient, the more similar the players' information, the more similar their assessments, and the less trade there is.

Key to our approach is to view the two players in our game as classical statisticians and tie their assessments to the actual outcome of the game. Instead of holding prior beliefs as in Bayesian models, it is as if the players had been engaged in a given outcome for a very long time. Each player was able to recover the distribution over her own payoffs from the interaction. In an equilibrium, the players then hold assessments which are consistent with such ideal observations and can be justified as having resulted from behavior of such consistent, rational and sophisticated individuals. Such assessments are called supporting assessments. Because the players' observations are partial and different, each player may hold a number of different supporting assessments. However, relative to Bayesian models with different priors, the equilibrium requirements in a given outcome impose more discipline on

supporting assessments. The welfare considerations from an individual's subjective perspective can then be based on any of her supporting assessments.⁵

The main difficulty with welfare assessments is that the individuals can hold a variety of statistically correct yet subjective assessments. The individuals can't tell what the true objective uncertainty is and it is no longer evident how the efficiency of an outcome should be evaluated. An outside observer would presumably face similar recovery problems unless she were magically endowed with much richer observations of the data than the individuals themselves. We tackle this issue by taking the perspective of each individual separately and considering her expected benefit from each outcome under any possible supporting assessment that she might hold. If an individual finds herself in a given equilibrium outcome and compares her expected gains with her expected gains she would have obtained in another equilibrium outcome. She finds the current equilibrium outcome preferable, if it is preferable under all supporting assessments that she might hold. Welfare evaluations are therefore intrinsically linked to equilibrium since players' assessments are determined in equilibrium.

This idea that one should take as a starting point the subjective welfare from each individual's perspective, to our knowledge dates at least as far as back as Wilson 1978. The idea is also motivated by Holmstrom and Myerson 1983, who define incentive-efficient allocations in Bayesian environments with a common prior. More recently, in a setting with heterogeneous priors which are not restricted by equilibrium considerations, Brunnermeier, Simsek, and Xiong 2013 develop a belief-neutral welfare measure. Their welfare measure can be thought of as a Pareto criterion based on each individual's subjective welfare criterion given all possible heterogeneous priors. They show that equivalently, their welfare measure can be thought of

⁵The individuals can be thought of as arriving to a specific assessments through a black box. Alternatively, precisely what assessment from amongst all her possible supporting assessments an individual holds is simply not known.

as a belief-neutral measure of social welfare from the perspective of a social planner.⁶ Here, players' assessments are determined in equilibrium, so that welfare must be measured differently.

In order to evaluate welfare from an individual's perspective, all possible assessment that she might hold in equilibrium may be considered. For example, each individual finds disagreement preferable to agreement under any possible supporting assessment that she might hold under disagreement. Therefore, in disagreement, the outcome subjectively Pareto dominates agreement. Ironically, agreement is fully recoverable, robust to the individuals' assessments over uncertainty, efficient, and, under the true distribution over uncertainty, it also maximizes each individual's expected gains. Under agreement, both individuals' assessments agree with the true objective uncertainty, and both find it preferable to disagreement – agreement subjectively (and objectively) Pareto dominates disagreement. Subjective Pareto evaluations are thus outcome dependent.

We have deliberately chosen to present our points in the form of an example that is as simple as possible. The set of possible states of the economy is discrete, there are only two representative individuals, each has only two possible actions, there are only two states of the world, and there are no dynamic considerations. The payoffs to the individuals are constructed in such a way that their equilibrium decision problems are essentially identical. The example given here is the simplest possible example of a disagreement that can only be supported by differing equilibrium assessments.⁷ Whereas the example is abstract to make our exposition transparent, our definitions are general.

⁶The motivation for their welfare measure comes partly from the literature in the decision theory pointing to the fact that under agents' conflicting beliefs a Pareto criterion might be problematic – see Mongin 1997, Gilboa, Samet, and Schmeidler 2004, and Gilboa, Samuelson, and Schmeidler 2012.

⁷In the language of Čopič 2014b, disagreement is pooling, informationally adverse, positive, and incentive imbalanced. By Theorem 1 and Proposition 6, each individual must have at least two possible signals and at least two possible actions for such an outcome to not be supportable under assessments that coincide.

In Section 2 we give the general form of our game, define the equilibrium and discuss notions of welfare when players' assessments are the same. In Section 3 specify the payoff parameters in our game, give an appropriate definition of welfare comparisons, and present our result. In Section 4 we relate the example to Aumann 1976. In Section 5 we relate the example to Milgrom and Stokey 1982. In Section 6, we recast the example in the context of autarky and exchange, and show how a central planner might be able to induce a switch from the inefficient autarky to the efficient exchange equilibrium in a way that balances the budget. In Section 7 we provide a general definition of welfare comparisons, and discuss welfare from the perspective of a social planner. Proofs are in the appendix.

2 The environment and the equilibrium definition

There are two players, $N = \{1, 2\}$, two states of the world, $\Omega = \{\omega_1, \omega_2\}$, each player can receive two signals, $\Theta_i = \{\theta_L, \theta_H\}$, and each player has two actions, $A_i = \{na_i, ag_i\}$. State ω_i is favorable to player i . The payoffs to player i in the two states and for different action profiles are given by a utility function $u_i : \Omega \times A \rightarrow R$. In our example, the payoffs to player i in the two states are specified by the following tables (i 's actions are in rows, and j 's actions are in columns):

ω_j	ag_j	na_j
ag_i	$1 - x^*$	$\underline{x}$
na_i	$\underline{y}$	y^*

ω_i	ag_j	na_j
ag_i	x^*	$\bar{x}$
na_i	$\bar{y}$	y^*

For example, when the state is ω_j , unfavorable to player i , player i plays ag_i and player j plays ag_j , the payoff to player i is $1 - x^*$. Of course, the payoff to player j is in that case x^* . We will later specify the relationships between the parameters $\underline{y}, \underline{x}, \bar{y}, \bar{x}, y^*, x^*$.

A player does not observe the state ω directly. Instead, she observes a signal θ_i . We assume that the aggregate signals to both players suffice to determine the state of the world with certainty: in ω_1 the signals to the players differ, $\omega_1 \equiv \{(\theta_L, \theta_H), (\theta_H, \theta_L)\}$; in ω_2 the signals to the players coincide, $\omega_2 \equiv \{(\theta_L, \theta_L), (\theta_H, \theta_H)\}$. State ω_1 is the *discordant* state, and ω_2 is the *concordant* state. All these facts regarding the payoff structure are common knowledge among the two players.

Since the aggregate signals to the two players completely determine the state, the strategic problem faced by the two players then has the following normal-form representation. Here the off-diagonal payoff matrices correspond to state ω_1 , favorable to player 1, and the diagonal payoff matrices correspond to state ω_2 , favorable to player 2.

(θ_L, θ_L)	ag_2	na_2
ag_1	$1 - x^*, x^*$	$\underline{x}, \bar{y}$
na_1	$\underline{y}, \bar{x}$	y^*, y^*

(θ_L, θ_H)	ag_2	na_2
ag_1	$x^*, 1 - x^*$	$\bar{x}, \underline{y}$
na_1	$\bar{y}, \underline{x}$	y^*, y^*

(θ_H, θ_L)	ag_2	na_2
ag_1	$x^*, 1 - x^*$	$\bar{x}, \underline{y}$
na_1	$\bar{y}, \underline{x}$	y^*, y^*

(θ_H, θ_H)	ag_2	na_2
ag_1	$1 - x^*, x^*$	$\underline{x}, \bar{y}$
na_1	$\underline{y}, \bar{x}$	y^*, y^*

What is not known to either player is the true (or objective) joint probability distribution $\bar{P}r$ over $\Omega \times_{i \in N} \Theta_i$, and the strategy of the other player. In our example, this objective probability distribution $\bar{P}r$ is specified by $\bar{P}r(\theta_i = \theta \mid \omega) = \frac{1}{2}$, $\forall \omega \in \Omega, \forall \theta_i \in \Theta_i$, $\bar{P}r(\theta_i = \theta_j \mid \omega_2) = \bar{P}r(\theta_i \neq \theta_j \mid \omega_1) = 1$, and $\bar{P}r(\omega_1) = \bar{P}r(\omega_2) = \frac{1}{2}$. Under the truth, each state is equally likely, and a signal to a player is by itself uninformative about the state or the signal to the other player.⁸

⁸Players know that the aggregate signals determine the state so that each player knows that certain draws of states and signals have probability zero, e.g., $\bar{P}r(\theta_L, \theta_L, \omega_1) = 0$. Players have no knowledge regarding the probability distribution over the draws that have non-zero probability.

A strategy s_i of player i is a contingent plan of action, $s_i : \Theta_i \rightarrow \Delta(A_i)$, where $s_i[\theta_i]$ denotes the probability distribution over i 's actions when her signal is θ_i ; and S_i denotes the set of i 's (mixed) strategies.⁹ An outcome is given by a pair $(\bar{P}r, s)$. An outcome realization is a draw of the state and signals $(\omega, \theta) \sim \bar{P}r$ and a draw of the action profile a , $a_i \sim s_i[\theta_i]$. We call the outcome where the strategy profile is given by (ag_1, ag_2) , *agreement*, and the outcome where the strategy profile is given by (na_1, na_2) , *disagreement*.

In the approach of classical statistics, the players do not have any priors over the uncertain parameters (including each other's strategy). Instead, the players form assessments based on their observations, where a player's observation can be imagined as an infinite dataset of independent realizations of a statistic of the outcome.¹⁰ An assessment by player i is a pair (Pr^i, s^i) , where $Pr^i \in \Delta(\Omega \times_{i \in N} \Theta_i)$, and $s^i \in S$. An assessment (Pr^i, s^i) is correct if $(Pr^i, s^i) = (\bar{P}r, s)$. Players' assessments need not be correct and are disciplined by probabilistic consistency with the parameters that a player can observe. A player observes her own signals, actions, and payoffs that she obtains. That is, in a given outcome $(\bar{P}r, s)$, a player observes the joint distribution over her own signals, actions, and payoffs. A player can thus verify whether her assessment is consistent with such a statistic.

The set of possible payoffs to player i is denoted by V_i , so that in our example, $V_i = \{\underline{y}, \underline{x}, \bar{y}, \bar{x}, y^*, 1 - x^*, x^*\}$. Given a (Pr, s) , denote by $\bar{P}r_{\Theta_i, A_i, V_i}[Pr, s]$ the induced

⁹When a player's strategy is deterministic and constant, i.e., independent of her signal, we denote it by the action that the player takes, e.g., ag_i denotes the strategy whereby player i takes the action ag_i regardless of her signal.

¹⁰The individuals are imagined as rigorous classical statisticians and decision makers who have collected infinite datasets of observations of equilibrium play. Their consistent assessments, equilibrium behavior, and justifications of others' behavior may be envisioned as stable points of an explicit dynamic process. See Čopić 2014a for such a dynamic justification of equilibrium points. Here, an explicit description of such a dynamic process would only unnecessarily complicate our formal exposition and make our points less clear.

joint probability over player i 's signals, actions, and payoffs.¹¹

DEFINITION 1. Player i 's assessment (Pr^i, s^i) is i -consistent with the outcome (Pr, s) if $\bar{Pr}_{\Theta_i, A_i, V_i}[Pr^i, s^i] \equiv \bar{Pr}_{\Theta_i, A_i, V_i}[Pr, s]$. The assessments are consistent when (Pr^i, s^i) is i -consistent with (Pr, s) , $\forall i \in N$.

For example, take the agreement outcome (ag_1, ag_2) , take player 1, and assume that all the payoff parameters are different and are also different from $1 - x^*$, in particular $x^* \neq \frac{1}{2}$. When she receives the signal θ_L , she takes the action ag_1 . She then observes her payoff x^* with the likelihood $\frac{1}{2}$, and that could only transpire if player 2 took the action ag_2 and the state was ω_1 so that player 2 must have observed the signal θ_L ; She observes her payoff $1 - x^*$ with the likelihood $\frac{1}{2}$ and that could only transpire when player 2 took the action ag_2 and the state was ω_2 so that player 2 must have observed the signal θ_H . Hence, player 1 deduces that, conditional on her signal, the states ω_1 and ω_2 occur with equal likelihoods, that player 2 observes her signals θ_L and θ_H with equal likelihoods, and that player 2 chooses ag_2 regardless of her signal. Similarly, when player 1 receives the signal θ_H she deduces that the states ω_1 and ω_2 occur with equal likelihoods, that player 2 observes her signals θ_L and θ_H with equal likelihoods and chooses ag_2 regardless. Since player 1 observes each of her signals with the likelihood $\frac{1}{2}$, it follows that she must correctly assess the other player's strategy and the likelihoods of each state and each signal. A similar argument is true for player 2. Therefore, in agreement, if a player holds a consistent assessment, her assessment is correct.

Our second definition concerns the players' incentive constraints. A player behaves optimally, given her assessment.

¹¹That is, for $\theta_i \in \Theta_i, a_i \in A_i, v_i \in V_i$,

$$\bar{Pr}_{\Theta_i, A_i, V_i}[Pr, s](\theta_i, a_i, v_i) = \sum_{\omega \in \Omega, \theta_j \in \Theta_j, a_j \in A_j | u_i(a_i, a_j, \omega) = v_i} Pr(\theta_j | \omega) s_j(a_j | \theta_j),$$

where $s_j(a_j | \theta_j)$ denotes the probability that j assigns to her action a_j when her signal is θ_j .

DEFINITION 2. Player i 's assessment (Pr^i, s^i) is i -incentive-compatible, i -IC, if:

$$\sum_{\theta_j \in \Theta_j, \omega \in \Omega} Pr^i(\omega \mid \theta_i, \theta_j) u_i(s_i^i(\theta_i), s_j^i(\theta_j); \omega) \geq \sum_{\theta_j \in \Theta_j, \omega \in \Omega} Pr^i(\omega \mid \theta_i, \theta_j) u_i(s_i'(\theta_i), s_j^i(\theta_j); \omega),$$

$$\forall \theta_i \in \Theta_i, \forall s_i' \in \Delta(S_i).$$

In a player equilibrium, a player optimizes given her assessment, she verifies that her assessment is consistent, and, assuming that her own assessment is the truth, she can impute such a consistent assessment on the other player – a player can justify the other player's behavior as a result of such an optimization and verification of consistency. A player equilibrium outcome thus satisfies three requirements: (i) optimization; (ii) consistency; (iii) justification. Additionally, requirement (iii) guarantees that even if the players had communicated their (possibly different) assessments to one another, neither one would have a reason to change her assessment.¹²

DEFINITION 3. An outcome (Pr, s) is a player equilibrium outcome if there exist assessments (Pr^1, s^1) and (Pr^2, s^2) such that,

1. (Pr^i, s^i) is i -IC, $i = 1, 2$,
2. (Pr^i, s^i) is i -consistent with (Pr, s) , $i = 1, 2$, and
3. (Pr^j, s^j) is j -consistent with (Pr^i, s^i) , $i, j = 1, 2$.

Assessments (Pr^i, s^i) , satisfying 1-3, are *supporting assessments*. Given a player equilibrium outcome (Pr, s) , denote by $O_i^{Pr, s}$ the set of all possible supporting as-

¹²Heuristically, player j can justify her observations as follows: given her assessment (Pr^i, s^i) , her observations can arise as a result of player j playing optimally under a consistent assessment, in particular, (Pr^j, s^j) ; and player j could then also justify her observations, and so on, ad infinitum. A formal epistemic characterization of player equilibrium is that the players' assessments are common belief, and satisfy optimization (or incentive constraints), consistency, a common belief in optimization and consistency (Čopić (2014a), Corollary 2). Player equilibrium outcomes can also be thought of as limit points of learning processes. Player equilibrium is related to Self-confirming equilibrium of Dekel et al (2004): if the justification requirement (iii) below is removed, then such an outcome is a Self-confirming equilibrium. Other related notions of equilibrium are Rationalizable Conjectural Equilibrium defined by Rubinstein and Wolinsky 1994, and Esponda 2013.

assessments for player i , $O_i^{Pr,s} = \{(Pr^i, s^i) \mid \exists (Pr^j, s^j) \text{ s.t. 1-3 hold}\}$.

An equilibrium outcome is *fully recoverable* if $|O_i^{Pr,s}| = 1, \forall i \in N$.

The standard notion of Bayes-Nash equilibrium can be imagined as a benchmark case where the players know $\bar{Pr}$ and each other's strategies so that their assessments are correct, whatever the reason may be. For example, that is true in a fully recoverable equilibrium outcome. A stronger notion of equilibrium is an *ex-post* Nash equilibrium where players need not know $\bar{Pr}$ – as long as player j plays her equilibrium strategy, player i 's equilibrium strategy is optimal, regardless of $\bar{Pr}$.¹³

DEFINITION 4. An outcome $(\bar{Pr}, s^*)$ is a Bayes-Nash equilibrium outcome, if,

$$\sum_{\theta_j \in \Theta_j, \omega \in \Omega} \bar{Pr}(\omega, \theta_j \mid \theta_i) u_i(s_i^*(\theta_i), s_j^*(\theta_j); \omega) \geq \sum_{\theta_j \in \Theta_j, \omega \in \Omega} \bar{Pr}(\omega, \theta_j \mid \theta_i) u_i(s'_i(\theta_i), s_j^*(\theta_j); \omega),$$

$$\forall \theta_i \in \Theta_i, \forall s'_i \in \Delta(S_i).$$

A strategy profile s^* is supportable in Bayes-Nash equilibrium, if there exists a Pr , such that (Pr, s^*) is a Bayes-Nash equilibrium outcome.

An outcome $(\bar{Pr}, s^*)$ is an *ex-post* equilibrium outcome if:

$$u_i(s_i^*(\theta_i), s_j^*(\theta_j); \omega) \geq u_i(s'_i(\theta_i), s_j^*(\theta_j); \omega),$$

$$\forall \omega \in \Omega \text{ and } \forall \theta_i \in \Theta_i, \theta_j \in \Theta_j, \text{ such that } \bar{Pr}(\theta_i, \theta_j \mid \omega) > 0, \text{ and } \forall s'_i \in \Delta(S_i).$$

A strategy profile s^* is supportable in *ex-post* equilibrium, if there exists a Pr , such that (Pr, s^*) is an *ex-post* equilibrium outcome.

By setting the players' assessments equal to $\bar{Pr}$, it is immediate that a Bayes-Nash equilibrium outcome is a player equilibrium outcome. In fact, a player equilibrium outcome is a Bayes-Nash equilibrium outcome, if and only if, the players' assessments

¹³Equivalently, rather than in expectation as in a Bayes-Nash equilibrium, player i 's strategy is a best reply to player j 's strategy point-wise, for every draw of uncertain parameters.

coincide.¹⁴ Additionally, when a Bayes-Nash equilibrium outcome is fully recoverable the players' assessments are correct. This is particularly relevant to our example here, so we state it as a proposition; we state the converse for completeness.

PROPOSITION 1. If an outcome is fully recoverable and is a Bayes-Nash equilibrium outcome, then the assessments are correct. Conversely, if an equilibrium outcome is supportable by assessments that are correct, then it is a Bayes-Nash equilibrium outcome.

When a Bayes-Nash equilibrium outcome (Pr, s) is fully recoverable, by Proposition 1, a player can evaluate welfare as she would in a Bayesian setting. The efficiency from an individual's perspective is then the *ex-ante* efficiency, similar to Wilson 1978. When an outcome is efficient for both players then it maximizes the *ex-ante* welfare. In a Bayesian game, there may be *ex-ante* efficient Bayes-Nash equilibrium outcomes in the sense of Holmstrom and Myerson 1983, i.e., outcomes that are Pareto-efficient in the *ex-ante* sense, which are not efficient for either player. The following definition of a welfare comparison, which applies to a fully-recoverable outcome $(\bar{Pr}, s)$ is a useful starting point for our welfare considerations of the next section.

DEFINITION 5. A fully-recoverable equilibrium outcome $(\bar{Pr}, s)$ is *i-efficient* relative to an equilibrium outcome $(\bar{Pr}, s')$ if:

$$\sum_{\theta \in \Theta, \omega \in \Omega} \bar{Pr}(\omega | \theta) u_i(s_i(\theta_i), s_j(\theta_j); \omega) \geq \sum_{\theta \in \Theta, \omega \in \Omega} \bar{Pr}(\omega | \theta) u_i(s'_i(\theta_i), s'_j(\theta_j); \omega).$$

If $(\bar{Pr}, s')$ is not an equilibrium outcome, then $(\bar{Pr}, s)$ is *i-efficient* relative to $(\bar{Pr}, s')$. The outcome $(\bar{Pr}, s)$ is *i-efficient* if $(\bar{Pr}, s)$ is *i-efficient* relative to $(\bar{Pr}, s')$, $\forall s' \in S$.

¹⁴Both of these implications are immediate; see Čopić 2014a, Corollary 1.

3 Information and welfare

In this section we give our example of welfare comparisons for which it is essential that one of the equilibrium outcomes is supportable under subjective and differing assessments. One equilibrium outcome may be subjectively Pareto superior to another equilibrium outcome and *vice-versa*. The first equilibrium outcome, disagreement, minimizes the sum of players' payoffs. The second equilibrium outcome, agreement, maximizes the sum of players' payoffs. Additionally, agreement is fully revealing and is an *ex-post* Nash equilibrium. The parameters in the above payoff structure are specified to satisfy these properties of the two equilibrium outcomes. The properties of agreement simplify the welfare comparison with disagreement. In this section we define the welfare comparison to the extent that is necessary to make our point; the example itself is as simple as possible to make that point.¹⁵

We now assume that parameters $\underline{y}, \underline{x}, \bar{y}, \bar{x}, y^*, x^*$ satisfy the following inequalities,

$$y^* < \frac{1}{2} < x^* < 1 \quad (1)$$

$$0 < \bar{x} - y^* = \alpha(y^* - \underline{x}), 1 < \alpha \quad (2)$$

$$\underline{y} < \bar{y} < 1 - x^* \quad (3)$$

$$2y^* < \bar{y} + \underline{x} < 1 \quad (4)$$

$$2y^* < \underline{y} + \bar{x} < 1. \quad (5)$$

Under these assumptions on parameter values, we have the following observations.

PROPOSITION 2. Under correct assessments $\bar{P}r$:

1. Agreement, (ag_1, ag_2) , is an *ex-post* equilibrium outcome and is fully recoverable.
2. Disagreement, (na_1, na_2) , is not supportable in Bayes-Nash equilibrium.

¹⁵We devote Section 7 to a general definition of welfare comparisons.

3. Agreement is i -efficient for $i \in N$ and maximizes the sum of players' utilities in every state.

Disagreement is not a Bayes-Nash equilibrium outcome, but it is an equilibrium outcome. It is not fully recoverable, so that there is for each player a set of possible supporting assessments. It is only supportable by assessments, which are different for the two players. We formally state this in the following proposition, where we denote disagreement by $(\bar{Pr}, \bar{s})$.

PROPOSITION 3. Disagreement $(\bar{Pr}, \bar{s})$ is a player equilibrium outcome. The supporting assessments (Pr^i, s^i) are such that $s^i \equiv \bar{s}$, and Pr^i are given by $Pr_{\Theta_j}^i(\theta_L) = \bar{Pr}_{\Theta_i}(\theta_L) = \frac{1}{2}$, $Pr_{\Theta_j}^i(\theta_H) = \bar{Pr}_{\Theta_i}(\theta_H) = \frac{1}{2}$, $i, j \in \{1, 2\}$, and,

$$\frac{Pr^1(\theta_L, \theta_L)}{Pr^1(\theta_L, \theta_H)} \geq \alpha \quad , \quad \frac{Pr^1(\theta_H, \theta_H)}{Pr^1(\theta_H, \theta_L)} \geq \alpha, \quad (6)$$

$$\frac{Pr^2(\theta_H, \theta_L)}{Pr^2(\theta_L, \theta_L)} \geq \alpha \quad , \quad \frac{Pr^2(\theta_L, \theta_H)}{Pr^2(\theta_H, \theta_H)} \geq \alpha. \quad (7)$$

By Proposition 3, the players' supporting assessments in disagreement must be different and that is common knowledge among the players. The set of each player's supporting assessments $O_i^{\bar{Pr}, \bar{s}}$ is given by:

$$O_1^{\bar{Pr}, \bar{s}} = \{Pr^1(\theta_L, \theta_L), Pr^1(\theta_H, \theta_H) \leq \frac{1}{1+\alpha}, Pr_{\Theta_1}^1(\theta_L) = Pr_{\Theta_1}^1(\theta_H) = \frac{1}{2}\}, \quad (8)$$

$$O_2^{\bar{Pr}, \bar{s}} = \{Pr^2(\theta_L, \theta_L), Pr^2(\theta_H, \theta_H) \geq \frac{\alpha}{1+\alpha}, Pr_{\Theta_2}^2(\theta_L) = Pr_{\Theta_2}^2(\theta_H) = \frac{1}{2}\}. \quad (9)$$

Suppose that a player is to compare her expected gain under disagreement to her expected gain under agreement. What should be assumed regarding the supporting assessment that she uses in order to make that comparison? Our view is that nothing should be assumed: any supporting assessment is as valid as any other supporting assessment – the player might in equilibrium hold any of these assessments, and she

might possibly even consider all such assessments.

Nevertheless, the players presumably cannot affect the underlying uncertainty. When comparing two equilibrium outcomes, both outcomes should be equilibrium outcomes under a player's assessment of underlying uncertainty. Thus, when comparing her gain under disagreement to her gain under agreement, a player in her assessment holds the underlying uncertainty as given. In agreement, both players' assessments regarding the strategy profile are correct since agreement is fully recoverable. In disagreement, both players' assessments regarding the strategy profile are also correct by Proposition 3. Since agreement is an *ex-post* Nash equilibrium it is an equilibrium under any assessment of the underlying uncertainty that supports disagreement. The following definition of subjective welfare comparisons therefore suffices for our example.

DEFINITION 6. Let (Pr, s) be an equilibrium outcome, such that $\tilde{s} \equiv s$, $\forall (\tilde{Pr}, \tilde{s}) \in O_i^{Pr, s}$, $\forall i \in N$. Let a strategy profile s' be supportable in a fully recoverable *ex-post* Nash equilibrium.

The outcome (Pr, s) is *i*-efficient relative to (Pr', s) , if (Pr', s) is *i*-efficient relative to (Pr', s') , $\forall (Pr', s) \in O_i^{Pr, s}$.

The outcome (Pr, s) subjectively Pareto dominates (Pr, s') if (Pr, s) is *i*-efficient relative to (Pr, s') , $\forall i \in N$.

Therefore, disagreement is *i*-efficient relative to agreement if it yields a higher expected gain under any assessment supporting disagreement in equilibrium. Since disagreement is an equilibrium outcome, i.e., incentive-feasible, under any such assessment, that is a valid comparison of the two outcomes. For a wide set of parameters, disagreement subjectively Pareto dominates agreement.

PROPOSITION 4. Suppose that $\alpha(1 - x^*) + x^* \leq (1 + \alpha)y^*$, where $\alpha = \frac{\bar{x} - y^*}{y^* - \bar{x}}$. Then disagreement subjectively Pareto dominates agreement.

Combining Proposition 2 with Proposition 3 yields an apparent contradiction. On the one hand, agreement is an equilibrium outcome, which is efficient under the objective uncertainty. In agreement, both players ascribe higher expected gains to agreement than to disagreement. On the other hand, disagreement is an equilibrium outcome, which is subjectively Pareto superior to agreement. In disagreement, each player ascribes a higher expected gain to disagreement under any subjective assessment that she might hold. The main point of our example is that both of these can simultaneously be valid. Heuristically, in disagreement players disagree regarding the redistributive properties of the equilibrium allocation due to their differing assessments over the underlying uncertainty, and we elaborate on this in the next section. Finally, we remark that for such a situation to arise it is necessary that at least one of the two outcomes be supportable only under assessments that are not correct and differ across the two players.¹⁶

4 Agreement and Disagreement

We now relate our example to Aumann's 1976 theorem.¹⁷ An equilibrium outcome can be imagined as one where players optimize, have consistent assessments, and each player's assessment can take the role of the truth so that the other individual's assessment must be consistent with that truth. Additionally, their supporting assessments have the property that, even if the players communicated their assessments to one another, no player would have any reason to change her assessment. In what follows, denote by $U_i^{Pr,s}$ the (subjective) expected payoff that individual i obtains under her assessment (Pr, s) .

When $\alpha(1 - x^*) + x^* < (1 + \alpha)y^*$, then, under any assessment that the other

¹⁶If disagreement were also supportable in Bayes-Nash equilibrium, then under such a common assessment one outcome would dominate the other or vice-versa, so that it would be impossible that the two outcomes mutually Pareto dominated one another.

¹⁷"If two people have the same priors, and their posteriors for an event A are common knowledge, then their posteriors are equal," Aumann 1976.

individual j might hold in disagreement, i would have preferred agreement over disagreement – this follows the argument in the proof of Proposition 4. We thus have the following relationships:

$$U_i^{Pr, \bar{s}} > U_i^{\bar{Pr}, s^*} > U_i^{Pr', \bar{s}}, \forall Pr \in O_i^{\bar{Pr}, \bar{s}}, \forall Pr' \in O_j^{\bar{Pr}, \bar{s}}, \forall i, j \in \{1, 2\}, j \neq i. \quad (10)$$

To interpret the relationships given by (10), consider player 1. Under any assessment that she may hold in disagreement, her subjective evaluation is that disagreement delivers superior expected gains than agreement; under any assessment that player 2 may hold in disagreement, player 1's expected gains would be higher in agreement than in disagreement. Therefore, in disagreement the assessments of player 1 imply that if the two players were to jointly change their strategies to agreement, player 2 would reap more than all the social gains from such a switch in regime. The relationships given by (10) similarly describe the welfare considerations of player 2.

In disagreement the two players disagree about the redistributive properties of equilibrium. This makes it impossible for the two players to jointly switch to the robust and efficient equilibrium outcome of agreement. Under any assessment that a player might hold, she is convinced that more than the total social gain from the switch in regime will go to the other player. Ironically, under the objective truth, agreement is preferred over disagreement also from the perspective of each player (Proposition 2). Since agreement is fully recoverable, were the individuals to switch to agreement, they would both agree about agreement being preferable.

Due to the relationships in (10), disagreement can be interpreted by the following narrative. Suppose player 2 were to suggest that the two individuals jointly deviate to agreement. Player 1 might then make the following argument: “I must conclude that you have realized that I have made the correct estimate and that my assessment is the truth. In that case, you will be better off in agreement. However, I will be

worse off and the only reason for you to suggest this switch to agreement is your self-interest, which is contrary to my interest.” Disagreement is therefore substantiated by a strong sense of subjective efficiency: in disagreement the two players disagree in their subjective assessments of uncertainty and both are convinced that disagreement is individually more efficient than agreement. By the third equilibrium requirement of justification, the players may maintain their disagreement even when their assessments are common knowledge.

This interpretation of players’ disagreement is complementary to Aumann 1976. By Aumann’s theorem, in a Bayesian setting, when two individuals hold a common prior and their posteriors are common knowledge, these posteriors must be the same. Here we can interpret the players’ assessments as priors, and their conditional assessments after receiving a signal as posteriors. The premises of Aumann’s theorem then hold under agreement. In contrast, in disagreement, players’ posteriors may be common knowledge, but since their priors do not coincide, their posteriors do not coincide either. It seems that such a statement might be easy to obtain in the context of priors of two individuals in a Bayesian setting, where these priors are not disciplined by consistency with equilibrium behavior. In the present setting of equilibrium and classical statistics the construction of appropriate assessments requires a careful argument involving equilibrium and welfare.

5 Trade and no-trade

We now reformulate our example to contrast it with another familiar theorem from the Bayesian setting. Assume (2) is changed into $x^* > y^* > 1 - x^* > 0$, while the other relationships (2)-(5) remain unchanged. This change does not affect any of the equilibrium considerations. In that case, the robust equilibrium outcome might yield a lower or higher aggregate welfare than the player equilibrium outcome, but the latter minimizes the risk. Both allocations are Pareto efficient, and we can

think of the robust equilibrium outcome as the initial allocation, and the player equilibrium outcome as the final allocation, after trade has taken place. Each player has traded-off some of her consumption in the high-payoff state for consumption in the low-payoff state, and has thus traded away some risk.¹⁸ We therefore interpret the robust equilibrium outcome (ag_1, ag_2) as “no-trade” and the player equilibrium outcome (nag_1, nag_2) as “trade.”

We can again interpret a player’s prior belief as her assessment, and her posterior belief as her conditional assessment after she has received her signal. Under no trade, the players’ assessments are correct and therefore coincide. The players can therefore be interpreted as having a common prior belief. Their posterior beliefs, i.e., their conditional assessments, are concordant. By Milgrom and Stokey 1982, there is no trade. In contrast, in the player equilibrium outcome players’ assessments are different, i.e., they do not have a common prior belief, and there is trade. Our formalization of the counterpoint between these two equilibrium statements in the context of classical statistics again requires a careful equilibrium construction.

Our interpretation here is intended as a comparative static between two different equilibrium outcomes. The players’ disparate assessments are a result of equilibrium behavior and can lead to trade where there would otherwise be none. In the sense that the “no-trade” outcome is fully recoverable, players then have more information than in the “trade” outcome. Hence, when players have more information, their assessments are necessarily more similar, and there is less trade. In the original version of our example one might be led to interpret disagreement as “less trade” and agreement as “more trade,” as the former is welfare minimizing, and the latter is welfare maximizing (see also next Section 6). Hence, in that version of the example, when players have more information, their assessments are more similar, and there is “more trade.” The present version is closer to the general-equilibrium environment

¹⁸If $y^* > \frac{1}{2}$ then the players can be thought of as risk averse, and such trade results in an ex-ante Pareto improvement.

described in Milgrom and Stokey (1982) in that the premises of the “no-trade” outcome correspond tightly to the premises in Milgrom and Stokey (1982). What these two versions of our example have in common is that when players have more information, in the sense that due to richer observations the outcome is (more) recoverable, their assessments are necessarily more similar.

6 Autarky and Exchange

Our example describes the trade-off between the idiosyncratic risks born by the players and the aggregate welfare. Under the third interpretation, we present our example as a toy macroeconomic example. If both players choose disagreement, which we here call *autarky*, then both receive the same payoffs in all states and there is no idiosyncratic risk. When both players choose agreement, or *exchange*, both can engage in a higher level of activity, but that brings some idiosyncratic risks to each player.

Now there are three possible states of the economy: an intermediate state and two extreme states – low and high. The discordant state with different signals is an “intermediate” state, $\omega_1 \equiv \omega_M$, and the concordant state ω_2 is now separated into two different states, $\omega_L \equiv (\theta_L, \theta_L)$ and $\omega_H \equiv (\theta_H, \theta_H)$, where ω_L is the “low” state, and ω_H is the “high” state. The action ag_i is interpreted as “exchange”, and $na g_i$ as “autarky,” i.e., agreement is now interpreted as exchange, and disagreement as autarky. We imagine the two extreme states as the low and the high points of the business cycle, and the intermediate state as the medium point of the business cycle. But that is for purely interpretative purposes: we assume that the players’ payoffs are specified by (1)-(5). Thus, when both individuals choose exchange, the aggregate welfare in the economy is maximal and is normalized to 1 – in the socially optimal outcome of exchange there is no aggregate risk.¹⁹

¹⁹While preserving all the necessary relationships for the equilibrium and welfare argument we

The relationships (1)-(5) now have the following interpretation. In exchange, the two individuals are ideal partners so that by (1), the total welfare in the economy – the sum of the players’ payoffs – is maximized when both individuals choose exchange. If an individual chooses exchange, we imagine that as a bet that one of her favorable states will transpire. Such a bet is on average (under the objective probability distribution over the states) more profitable than remaining in autarky, even if the other individual chooses autarky – by (2). By (3), if an individual chooses autarky while the other chooses exchange, then the autarchic individual is penalized in a state that is unfavorable to her, e.g., she might lose some of her market share to competitors. In a favorable state, the externality on the autarchic individual is positive, e.g., her competitors’ relative market shares have been diminished. When both individuals choose autarky, neither of them faces any externalities or idiosyncratic risks, and by (4) and (5), that minimizes the aggregate welfare.

In this version, we can imagine the two players as representative individuals of two different sectors. Player 1 is a representative individual of a productive sector and player 2 is a representative individual of a financial, or speculative sector. For player 1, autarky bears lower risks but also results in a lower level of production; exchange bears higher risks and results in a higher average production. For example, exchange can be interpreted as taking a loan collateralized by the output of her production, which enables her to produce at a higher level. In the intermediate state of the economy, her bet pays off so that she can fully repay the loan and obtain a higher profit than had she remained in autarky. In an extreme state of the economy (low or high), she cannot fully repay her loan, so that her output is seized, which results in lower gains relative to autarky. In exchange, she is therefore facing some idiosyncratic risks that she might, to a certain extent, be able to trade away by buying

could also assume that in exchange, the aggregate welfare is highest in ω_H , intermediate in ω_M , and lowest in ω_L , and then normalize the average welfare in exchange to 1. For example, that could be achieved by multiplying the payoffs in exchange and the deviations therefrom in ω_H by a factor $(1 + \beta)$, and in ω_L by $(1 - \beta)$, for some $\beta \in (0, 1)$, such that $x^*(1 - \beta) > y^*$.

insurance against the unfavorable states.²⁰ For the speculative individual, autarky can be imagined as a low-risk investment in government bonds, while exchange can be imagined as a risky investment in the productive sector. However, the speculative individual makes gains precisely by acquiring the distressed assets of the productive sector in the extreme states of the economy, and makes losses relative to autarky when the state of the economy is intermediate.

We recall our results of Section 3. Exchange is a fully-recoverable robust equilibrium outcome, which maximizes aggregate welfare. Autarky is an equilibrium outcome, which is only supportable under differing assessments. Exchange Pareto dominates autarky, and autarky subjectively Pareto dominates exchange. Therefore, if the individuals encounter themselves in autarky, neither would be willing to switch to exchange. Since exchange maximizes aggregate welfare, it seems there might be a role for subsidies to induce the individuals to switch from autarky to exchange.²¹ An individual would be willing to switch to exchange in any state if the subsidy provided sufficient compensation under her assessment that was most favorable to autarky. An assessment is most favorable to autarky when all the probability mass is assigned to the unfavorable states – such an assessment is an equilibrium assessment by (8) and (9). The following Proposition 5 states the amount of such subsidy. If the individuals were “less pessimistic,” in the sense that their assessments were bounded away from the assessments most favorable to autarky, then smaller inducements might suffice. However, in general that cannot be known to the planner.

²⁰Non-luxury cars fit this description well. In a low state of the economy, the demand for cars is low. In a high state of the economy, the demand for cars may be high but the consumers may instead buy luxury cars, so that the demand for non-luxury cars may nonetheless be low. More generally, production of a good for which there exists a substitute giffen good may correspond well to the productive sector described here.

²¹We assume there is a social planner, who observes, for example, the total sum of the individuals’ expected payoffs in this stylized economy. That is enough to determine whether the individuals are in autarky or in exchange. If the planner observed nothing at all about the interaction, then the planner’s problem would make no sense. We could also assume that the planner observed the distribution over the individuals’ actions, as in Čopić (2014b), and in the present example that would make no difference.

PROPOSITION 5. Suppose $\alpha(1 - x^*) + x^* < (1 + \alpha)y^*$. Then, for the individuals to switch from autarky to exchange under any supporting assessment, each must be given a subsidy in the amount of $y^* + x^* - 1$.

In the present example there is no aggregate risk under exchange. The subsidies to the individuals can thus be made contingent on the individuals' payoffs under exchange, which would balance the budget. For example, a planner could promise each individual to compensate her in case of a loss and in turn take some fraction β of the individual's extra earnings in case of a gain. Since exchange is here efficient while autarky is not, it follows that for the planner to balance the budget, $\beta < 1$, so that both individuals should presumably be eager to take such a bet. The planner could thus induce the change from the expectations-driven inefficient autarky to the efficient exchange by providing a subsidy, which she could later recoup.

Such policy interventions presuppose that the planner were able to observe and enforce the change in the individuals' behavior. Otherwise, both individuals might collect their subsidies and remain in autarky. If the planner is unable to recoup the subsidy, or the amount of subsidy is relevant in the short run, before it can be recovered, then the magnitude of idiosyncratic risk in the efficient outcome matters. That is, the lower the difference $x^* - \frac{1}{2}$, the lower the magnitude of idiosyncratic risk and the lower the amount of subsidy $x^* + y^* - 1$. Note that there can be some idiosyncratic risk for the necessary subsidy to equal zero – by Proposition 4, when $x^* < \frac{\alpha - (1 + \alpha)}{\alpha - 1}$, autarky is no longer subjectively efficient.

Everything else equal, what determines the magnitude of idiosyncratic risk is how much of it can be traded away. Here we do not explicitly model the manner in which the idiosyncratic risk is traded away. Instead, we revisit our interpretation of player 1 as the productive sector and player 2 as the financial sector of the economy. The way in which player 1 engages in exchange is by taking a collateralized loan, where her production serves as the collateral: ω_L is the state in which her realized production

is low, and ω_H is the state in which the demand for her product is low – under the interpretation that ω_H is the high state of the economy, that can be the case for example if for the good produced by player 1, there is a substitute commodity, which is a giffen good. In the intermediate state ω_M , player 1 makes profits. For the financial sector, such collateralized loan to the productive sector is on average more profitable than a safer investment under autarky, e.g., government bonds. Therefore, player 1 may be able to trade away the idiosyncratic risk by insuring against the states ω_L and ω_H . For example, she might go short some number of call options on her production at a high strike price (to insure against ω_H), and go long some number of put options at a low strike price (to insure against ω_L). In what amount player 1 can insure against the extreme states of the economy thus depends on the availability and prices of such possible insurance schemes. That is, it depends on the counterpart of that trade taken up by the financial sector. The sort of examples that fit the present model are those where the idiosyncratic risk cannot be traded away entirely. We thus consider it as an assumption that in the economy studied here, some residual idiosyncratic risk is a property of the exchange equilibrium outcome.

Nevertheless, in some cases, where the reasons for residual idiosyncratic risks are of institutional nature, there may exist remedies for the problems described here. In this example, if the residual idiosyncratic risks is a result of poor access to financial markets by the productive sector, then improving the ability of the productive sector to trade away its risks would remedy the problem described here. Alternatively, if such residual idiosyncratic risk stemmed from issues with pricing, e.g., a collusive scheme or monopolistic power in one of the sectors of the economy, then ensuring more competition in that sector would likely resolve the problem. In the most favorable situation, all the idiosyncratic risk in the efficient and robust equilibrium of exchange can be traded away. In that case $x^* = \frac{1}{2}$ and there no longer exist any assessments under which autarky would be subjectively efficient; it is also no longer a player equilibrium outcome. Once there are no limits to trading away

idiosyncratic risks, then exchange is a unique equilibrium outcome, which is fully recoverable, robust, and maximizes individual and social welfare in every state.

7 General welfare comparisons

Definition 6 is suited to our example and circumvents some of the problems concerning welfare comparisons of two different outcomes. In this section we define welfare comparisons for the general case. The efficiency comparison necessarily embodies both, the supporting assessments, as well as incentive constraints, which yield restrictions on these supporting assessments. The difference between Definition 6 and the following general Definition 7 is that in the latter, there is no assumption on the properties of the outcome under comparison with the equilibrium outcome. The simplest way to do that is by defining an equilibrium outcome to be more efficient than an outcome, which is not an equilibrium outcome under the given assessment – a non-equilibrium outcome does not satisfy incentive feasibility.

Key to the definition is that in different equilibrium outcomes players may hold multiple, and possibly different, supporting assessments. Such multiplicity is relevant to the welfare comparison of these outcomes. Suppose in an equilibrium outcome a player holds an assessment (Pr, s) . She can then consider her expected gain in the current outcome relative to some other outcome (Pr, s') . Since players presumably have no effect on the objective uncertainty, everything that can change in the alternative outcome are the strategies that the players play. For the comparison to be viable, the outcome (Pr, s') must be an equilibrium outcome. If (Pr, s') were not an equilibrium outcome under her current assessment of the underlying uncertainty, then such a comparison would be meaningless – even if the alternative outcome delivered a higher expected gain, that outcome would not be viable in terms of incentive feasibility. For that reason, in our general definition, an equilibrium outcome is defined to be more efficient than an outcome, which is not an equilibrium outcome.

A player's assessment of the strategy profile is subjective. When making a comparison between (Pr, s) and (Pr, s') the player may compare the expected gain under any assessment $(\tilde{Pr}, \tilde{s}) \in O_i^{Pr, s}$ with the expected gain under $(\tilde{Pr}, \tilde{s}') \in O^{Pr, s'}$. The following definition of efficiency and Pareto criterion extends Definition 6 to the general case.

DEFINITION 7. If (Pr, s') is not an equilibrium outcome, then an equilibrium outcome (Pr, s) is *i-efficient* relative to (Pr, s') .

An equilibrium outcome (Pr, s) is *i-efficient* relative to an equilibrium outcome (Pr, s') , if (Pr^i, s^i) is *i-efficient* relative to $(Pr^i, \tilde{s}^i)$, $\forall Pr^i, s^i, \tilde{s}^i$, s.t., $(Pr^i, s^i) \in O_i^{Pr, s}$ and $(Pr^i, \tilde{s}^i) \in O_i^{Pr, s'}$;

The outcome (Pr, s) subjectively Pareto dominates (Pr, s') , if the outcome (Pr, s) is *i-efficient* relative to (Pr, s') , for all $i \in N$.

In the presence of a social planner, a different welfare comparison can be envisioned. In Section 6, the social planner observed some statistic of the interaction between the two players – the sum of their expected payoffs from the interaction. Exchange was socially superior to autarky because it generated a superior sum of payoffs to the players under any assessment under which payoffs consistent with autarky were possible in equilibrium. That would suggest an *objective* utilitarian social welfare criterion (with equal weights on the players' payoffs), where an outcome is socially superior if it generates higher expected gains under any supporting assessment. Such a welfare criterion results in a partial ordering of equilibrium outcomes. The example of Section 6 illustrates how a social planner can induce the change to a socially superior outcome through appropriate subsidies, which can in that example be recouped. More importantly, it illustrates that such a utilitarian welfare criterion does not coincide with the subjective Pareto criterion: autarky is subjectively Pareto superior to exchange, while exchange is socially superior to autarky.

8 Appendix

Proof of Proposition 1. Let (Pr, s) be a Bayes-Nash equilibrium outcome, so that (Pr, s) satisfies i -IC, for all $i \in N$. Additionally, (Pr, s) is i -consistent with (Pr, s) , and since $|O_i^{Pr, s}| = 1$, $O_i^{Pr, s} \equiv \{(Pr, s)\}$, $\forall i \in N$. The converse follows from Definition 4. $\square$

Proof of Proposition 2. Statement 1 is immediate to verify. To see 2, take player 1, and her type θ_H . If she deviates from na_1 , then when $\theta_2 = \theta_L$ her payoff increases by $\bar{x} - y^*$, and when $\theta_2 = \theta_H$ her payoff decreases by $y^* - \underline{x}$. Since $\bar{x} - y^* = \alpha(y^* - \underline{x})$, where $\alpha > 1$, it follows that under $\bar{Pr}$ she would have incentives to deviate, so that disagreement is not a Bayes-Nash equilibrium outcome. Point 3 is also immediate to verify. $\square$

Proof of Proposition 3. First note that in disagreement, in order for i -consistency to hold, player i must correctly assess the other player's strategy, i.e., $s^i \equiv \bar{s}$; she must also correctly assess the likelihoods of her own signals, i.e., $Pr_{\Theta_i}^i(\theta_i) = \frac{1}{2}$, $\theta_i \in \{\theta_L, \theta_H\}$; but her payoffs do not vary with signals to the other player, so that for consistency with $\bar{Pr}$, conditional on her own signal, she can make any assessment regarding the likelihoods of the signals to the other individual. In order for the equilibrium requirement (iii) to hold, that is, for Pr^j to be j -consistent with Pr^i , it must be that $Pr_{\Theta_j}^i(\theta_j) = \frac{1}{2}$, $\theta_j \in \{\theta_L, \theta_H\}$.

In disagreement, the incentive constraints of player 1 are given by,

$$Pr^1(\theta_L, \theta_L)(\underline{x} - y^*) + Pr^1(\theta_L, \theta_H)(\bar{x} - y^*) \leq 0,$$

$$Pr^1(\theta_H, \theta_L)(\bar{x} - y^*) + Pr^1(\theta_H, \theta_H)(\underline{x} - y^*) \leq 0,$$

and the incentive constraints of player 2 are given by,

$$Pr^2(\theta_L, \theta_L)(\bar{x} - y^*) + Pr^2(\theta_H, \theta_L)(\underline{x} - y^*) \leq 0,$$

$$Pr^2(\theta_L, \theta_H)(\underline{x} - y^*) + Pr^2(\theta_H, \theta_H)(\bar{x} - y^*) \leq 0.$$

Since $(\bar{x} - y^*) = \alpha(\underline{x} - y^*)$, the claim follows. $\square$

Proof of Proposition 4. Take player 1 and suppose her signal is θ_H . In disagreement, $(\bar{P}r, \bar{s})$, any of her supporting assessments must put a sufficiently large probability mass on the draw of signals (θ_H, θ_L) . From Proposition 3,

$$P^i(\theta_H, \theta_L) \geq \frac{1}{2} \frac{\alpha}{1 + \alpha}.$$

The expected payoff to the player in disagreement is y^* . Since $x^* > 1 - x^*$, in agreement her subjective expected payoff (conditional on θ_H) is decreasing in the probability mass that she assigns to the event (θ_H, θ_L) . Therefore, her subjective expected payoff in agreement is highest when $P^i(\theta_H, \theta_L) = \frac{1}{2} \times \frac{\alpha}{1 + \alpha}$. By a similar argument when she observes the signal θ_L , we conclude that in disagreement, the highest subjective expected payoff that she can assign to agreement equals,

$$\frac{\alpha}{1 + \alpha}(1 - x^*) + \frac{1}{1 + \alpha}x^*.$$

Thus, her subjective expected payoff will be higher in disagreement than in agreement under any supporting assessment if,

$$\frac{\alpha}{1 + \alpha}(1 - x^*) + \frac{1}{1 + \alpha}x^* \leq y^*.$$

A similar argument can be applied to Individual 2. $\square$

Proof of Proposition 5. When player i assigns all the probability mass to the unfa-

vorable state, such an assessment supports autarky in equilibrium. The player then obtains y^* under autarky and $1 - x^*$ under exchange, so that the necessary subsidy for her to switch to exchange is $y^* - (1 - x^*) = y^* + x^* - 1$. Player i 's payoff under exchange is weakly greater than $1 - x^*$ under any assessment supporting autarky. $\square$

References

- [1] AMADOR, M. AND P. O. WEILL [2010]: “Learning from Prices: Public Communication and Welfare,” *Journal of Political Economy*, **118**, 866 - 907.
- [2] ANDERSON, R. M. AND H. SONNENSCHN [1985]: “Rational Expectations Equilibrium with Econometric Models,” *Review of Economic Studies*, **52**, 359 - 369.
- [3] ANGELETOS, G. M. AND J. LA’O [2013]: “Sentiments,” *Econometrica*, **81**, 739 - 779.
- [4] AUMANN, R. [1976]: “Agreeing to Disagree,” *The Annals of Statistics*, **4**, 1236-1239.
- [5] BERGEMANN, D. AND S. MORRIS [2005]: “Robust Mechanism Design,” *Econometrica*, **73**, 1771-1813.
- [6] ČOPIČ, J. [2012]: “Nash equilibrium, rational expectations, and heterogeneous beliefs: Action-consistent Nash equilibrium,” mimeo.
- [7] ČOPIČ, J. [2014a]: ‘Normal-form Games with Decentralized Rational Statisticians: Moral hazard and Adverse selection,’ mimeo.
- [8] ČOPIČ, J. [2014b]: “Players-Econometricians, Adverse Selection, and Equilibrium,” mimeo.

- [9] DEKEL E., D. FUDENBERG, AND D. K. LEVINE [2004]: “Learning to Play Bayesian Games,” *Games and Economic Behavior*, **46**, 282 - 303.
- [10] ESPONDA, I. [2013]: “Rationalizable Conjectural Equilibrium: A Framework for Robust Predictions,” *Theoretical Economics*, *forthcoming*.
- [11] FOSTEL, A. AND J. GEANAKOPOLOS [2008]: “Leverage cycles and the anxious economy,” *American Economic Review* **98**, 1211-1244.
- [12] GEANAKOPOLOS, J. [2010]: “The Leverage Cycle,” in *NBER Macroeconomics Annual*, ed. by D. Acemoglu, K. Rogoff, and M. Woodford. Cambridge, MA: MIT Press, 1 - 65.
- [13] GILBOA, I., D. SAMET, AND D. SCHMEIDLER [2004]: “Utilitarian aggregation of beliefs and tastes,” *Journal of Political Economy* **112**, 932-938.
- [14] GILBOA, I., L. SAMUELSON, AND D. SCHMEIDLER [2012]: “No-betting Pareto dominance,” *mimeo*.
- [15] HARRISON, M. AND D. KREPS [1978]: “Speculative investor behavior in a stock market with heterogeneous expectations,” *Quarterly Journal of Economics* **92**, 323-336.
- [16] HOLMSTROM, B. AND MYERSON R. B. [1983]: “Efficient and Durable Decision Rules with Incomplete Information,” **51**, 1799 - 1820.
- [17] KURZ, M. [2011]: “Symposium: on the role of market belief in economic dynamics, an introduction,” *Economic Theory*, **47**, 189-204.
- [18] MILGROM P., AND N. STOKEY [1982]: “Information, Trade, and Common Knowledge,” *Journal of Economic Theory*, **26**: 17-27.
- [19] MILLER, E. [1977]: “Risk, uncertainty, and divergence of opinion,” *Journal of Finance* **32**, 1151-1168.

- [20] MONGIN, P. [1997]: “Spurious unanimity and the Pareto principle,” mimeo.
- [21] MORRIS, S. [1995]: “The Common-Prior Assumption in Economic Theory,” *Economics and Philosophy*, **11**, 227 - 253.
- [22] ORDOÑEZ, G. [2013]: “The Asymmetric Effects of Financial Frictions,” *Journal of Political Economy*, **121**, 844 - 895.
- [23] RADNER R. [1979]: “Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices,” *Econometrica*, **47**, 655-678.
- [24] RADNER, R. [1982]: “Equilibrium under uncertainty,” Chap. 20 in *Handbook of Mathematical Economics*, vol. II, edited by K. Arrow and M. D. Intrilligator. Amsterdam: North-Holland.
- [25] RUBINSTEIN, A. AND A. WOLINSKY [1994]: “Rationalizable Conjectural Equilibrium: Between Nash and Rationalizability,” *Games and Economic Behavior*, **6**, 299-311.
- [26] SCHEINKMAN, J. AND W. XIONG [2003]: “Overconfidence and speculative bubbles,” *Journal of Political Economy*, **111**, 1183-1219.
- [27] SIMSEK, A. [2013]: “Belief Disagreement and Collateral Constraints,” *Econometrica*, **81**, 1-53.
- [28] WILSON, R. [1978]: “Information, Efficiency, and the Core of an Economy,” *Econometrica*, **46**, 807-816.