Decentralized trading with private information*

Michael Golosov
MIT and NES

Guido Lorenzoni
MIT

Aleh Tsyvinski
Harvard and NES

April 2008

Preliminary draft. Do not circulate without permission.

Abstract

We contribute to the recently developed theory of asset pricing in decentralized markets. We extend this literature to characterize an environment in which some agents have superior private information. In our model, agents have an additional incentive to trade assets to learn information that other agents have. First, we show that uninformed agents can learn all the useful information the long run, and that the long-run allocations are Pareto efficient. In the long run, therefore, the allocations coincide with those of the standard centralized market equilibrium such as in Grossman-Stiglitz. Second, we show that agents with private information receive rents, and the value of information is positive. This is in contrast with the centralized markets in which prices fully reveal information and the value of information is zero. Finally, we provide characterization of the dynamics of the trades.

1 Introduction

This paper provides a theory of trading and information in environments which are *informationally decentralized*. These markets have three key frictions: (1) trading is decentralized (bilateral), (2) information about transactions is known only to the parties of the transaction, and (3) some agents have private information. Duffie, Garleanu, and Pedersen (2005) started a research agenda of providing a theory of asset pricing in decentralized environments with public information.[1] They note that many important markets are decentralized such as over-the-counter markets and private-auction markets. Examples of such markets include mortgage-backed securities, swaps and many other derivatives, and real estate markets to name a few. Many of these markets feature informational frictions as well which are

*We thank Daron Acemoglu, Marios Angeletos, Gady Barlevy, Marco Bassetto, V.V. Chari, Darrell Duffie, John Geanakoplos, and Ivan Werning and seminar audiences at Columbia, Chicago Fed, HEC Lausanne, NYU, and Yale for useful comments.

[1]There is a large literature now studying such markets – see, e.g., Duffie, Garleanu and Pedersen 2007, Duffie and Manso 2007, Lagos 2007, Lagos and Rochetau 2007, Lagos, Rochetau, and Weill 2007, Vayanos 1998, Vayanos and Weill 2007, and Weill 2007, among others.

a focus of our paper – prices for transactions are not publicly observable and some agents are better informed than others.

We are motivated by three interconnected sets of issues. The first is whether trading in informationally decentralized markets leads to an efficient outcome. This is an important open issue because the presence of any of the three key frictions may lead to highly inefficient outcomes. Our second focus is on the value of information and the evolution of such value in our environment. Are the informed agents better off than uninformed? Can and should the uninformed agents learn private information? This is one of the classis issues of the asset pricing literature in environments with private information and *centralized* tradingsuch as Grossman and Stiglitz (1980). They conclude that prices are fully revealing (in the absence of noise traders), and private information has no value. Our analysis answers these questions in our decentralized environment. Finally, our interest is in the dynamics of trades. In this regard, we are motivated by the classic analysis of Glosten and Milgrom (1985) and Kyle (1985). The difference with this literature is that there markets are informationally centralized in the sense that all agents observe all transactions. In contrast, in our environment information about transactions is private.

Specifically, our environment is as follows. Agents start with different endowments of two assets, match randomly, and trade in bilateral meetings. Any information about their trade is private to the parties of the transaction. A proportion of agents are informed and have superior information over the uninformed agents. This information is about the probability that an asset will pay off in a given state of the world and determines how valuable the asset is. In each period, the game can end with some probability, and the agents have to consume their endowments of assets, or the game continues to the next period. This formulation is one way to introduce discounting in the model. The only information that the agents observe are the history of their matches, but not the endowments of other agents or their trades. Uninformed agents form beliefs about the value of an asset based on a history of trades they conduct. This environment is technically and conceptually challenging to analyze because the distribution of beliefs about the value of the asset is endogenous and changes over time. An uninformed agent not only has to form a belief about the state of the world but also to form a belief about other agents' beliefs as they influence the future opportunities of trading.

We derive two sets of results. The first set of results are a theoretical examination of efficiency of equilibrium and of the value of information and its evolution. We first show that the long-run allocations are Pareto efficient, and our decentralized environment converges to allocations achieved in Grossman-Stiglitz' perfectly revealing equilibrium. The argument is by contradiction. If an uniformed agent does not converge to an efficient allocations there is a profitable deviation on his part where he constructs a trade thal allows him to learn the state of the world and then take advantage of this information. We show that the losses of such experimentation can be made smaller than the gains of learning the state of the world.

If the initial allocations are not Pareto-efficient, i.e., if there are gains from trade[2], the informed

[2] If the initial allocations are already Pareto optimal, we show that a version of no trade theorem (similar to, e.g., Brunnermeier 2001 for a detailed exposition of this topic) holds.

agents receive a higher lifetime utility than uninformed agents. In other words, private information has a positive value. The intuition is that the uninformed agents will learn the true state of the world only in the long run and additionally have to conduct potentially unprofitable trades in the short run to learn the state of the world. That is why in the short run there are profitable trading opportunities for the informed agents. This result is in contrast with the Grossman-Stiglitz analysis where prices fully reveal all the private information, and information has no value. Yet, the uninformed agents can learn all the usefull information, and, in the long run, the value of information converges to zero.

The second set of results is on the theoretical and numerical analysis of the dynamics of trades. We first consider a static example in which there is only one round of trading that is useful to illustrate the inutiition about the trades and strategies of the agents. We show that the static allocations are inefficient. We then develop a method to numerically compute a specific equilibrium of the game. Our simulations and examples show how the behavior of informed agents differ depending on their endowment of the valuable asset. The asset position of the agent who starts with a low endowment of the valuable asset follows a hump-shaped profile. This agent accumulates the valuable asset above his long-run position before the information is revealed. To do so, he mimics the behavior of the uninformed agents and takes advantage of the fact that uninformed agents do not know which asset is more valuable. Upon accumulating a sufficient amount, this agent sells some of the valuable asset at more advantagous terms of trade as information dissipates across agents. The strategy of the informed agent with a large initial endowment of the valuable asset is different. He decumulates his endowment of the valuable asset. His strategy is determined by considerations of signalling that his asset is valuable – to do so he exchanges small amounts of assets for large amounts of the other asset. Finally, we show in the examples that it takes longer to converge to efficient allocations in our environment with private information than if all information is public.

Our paper is related to several other strands of the literature. The most closely related is Duffie and Manso (2007) and Duffie, Giroux, and Manso (2007) who also consider a private information trading setup with decentralized markets and focus on information percolation in these environments. They derive important closed form solutions for the dynamics of the trade in an environment similar to ours while we have a more general setup and derive strong results about the long-run allocations and general dynamics. Amador and Weill (2007, 2008) is an interesting study of information dispersion in an environments with private and public information. The difference in this paper is that ours is a model of trade rather than solely of information transmission.

Our work is also related and extends papers by Wolinsky (1990) and Blouin and Serrano (2001) who consider a version of Gale (1987) economy with indivisible good and heterogenous information about its value. They show that the information is not fully revealed and allocations are not ex-post efficient. The difference of our paper is that we allow for endogenously determined prices rather than assuming fixed terms of trade. Dubey, Geanakoplos, and Shubik (1987) and Glosten and Milgrom (1985) is related but they consider a model where there are commonly observed signals (”prices”) through which uninformed agents learn. In our environment all prices are determined as a part of equilibrium.

The paper is structured as follows. Section 2 describes the environment. Section 3 defines an equilibrium of the game. Section 4 provides characterization of the equilibrium. Section 5 is a static example. Section 6 is a numerical solution of the game. Section 7 concludes. The Appendix contains most of the formal proofs which are sketched in the body of the paper.

2 Setup and trading game

This section describes the setup of our model and defines the decentralized trading game.

2.1 Environment

There are two states of the world $S \in \{S_1, S_2\}$ and two assets. The asset $j \in \{1, 2\}$ pays one unit of consumption if and only if the state S_j is realized. There is a continuum of agents divided in a finite number of types N of the allocations of the initial endowment. Each type $i \in N$ has an initial endowment of the two assets, denoted by the vector $x_{i,0} \equiv \left(x^1_{i,0}, x^2_{i,0}\right)$, where the second subscript indicates the time period. All agents have identical von-Neumann-Morgenstern expected utility $E\left[u(c)\right]$, where E is the expectation operator. Let f_i denote the fraction of agents of type i. We normalize the total initial endowment of assets to 1 in each state:

$$\sum_i f_i x^j_{i,0} = 1 \text{ for } j = 1, 2. \tag{1}$$

We make the following assumptions on preferences and endowments. The first assumption is symmetry insuring that the endowments of assets are mirror images of each other.

Assumption 1. *(Symmetry) For each type $i \in N$ there exists a type $j \in N$ such that $f_i = f_j$ and $(x^1_{i,0}, x^2_{i,0}) = (x^2_{j,0}, x^1_{j,0})$.*

The second assumption imposes the usual properties of the utility function as well as boundary and Inada conditions.

Assumption 2. *The utility function $u(\cdot)$ is differentiable on R^2_{++},* continuous, increasing, strictly concave, *bounded above, and satisfies $\lim_{x\to 0} u(x) = -\infty$.*

Finally, we assume that the initial endowments are interior.

Assumption 3 *The initial endowment $\left(x^1_{i,0}, x^2_{i,0}\right)$ for all agents i is in the interior of R^2_+.*

The uncertainty about the state of the world is realized in two stages. First, nature draws a binary signal $s \in \{s_1, s_2\}$, with equal probabilities. Second, given the signal s, nature chooses state S_1 with probability $\phi \in (1/2, 1)$ if $s = s_1$ and with probability $1 - \phi$ if $s = s_2$. After the signal s is realized, an exogenous random fraction α of agents in each group privately observes the realization of s. The agents who observe s are the *informed* agents, denoted by I, the ones who do not observe s are the *uninformed* agents, denoted by U. In other words, the informed agents know the probability of the true state of the world, while the uninformed agents have a prior equal to 0.5. Throughout the paper, we assume that the agents do not observe the information set of their counterparts, and that endowments are privately observable.

2.2 Trading

After the realization of the signal s, but before the realization of the true state S, all agents engage in a trading game. This game is set in discrete time.

In each period t, the trading game can either end or continue. At the beginning of each round of trading t, with probability $1-\gamma$ the game ends, the state S is publicly revealed, and agents consume the payoffs of their assets. The possibility that the game ends is one reason for agents to trade assets as they want to insure themselves. If the game does not end, all agents are randomly matched in pairs, hence, this game is a decentralized, bilateral trading setup.

The trading within the match happens as follows. With a probability 0.5 one of the agents is selected to propose a take it or leave it offer $z=(z_1,z_2)\in R^2$ to the other agent in the match, i.e., offers to trade z_1 of the first asset for z_2 of the second asset. The other agent in the match can either accept or reject the offer. If an agent with endowment x offers z to an agent with $\tilde{x}$ and the offer is accepted, the endowment of the agent who made an offer becomes $x-z$, and the endowment of the agent who accepted the offer becomes $\tilde{x}+z$. The offer must be feasible: $x-z\geq 0$, and can be accepted only if $\tilde{x}+z\geq 0$. If the offer is rejected, both agents remain with the same endowments as at the start of the bargaining round. This concludes the bargaining round t. Notice that, except for the presence of asymmetric information, this game follows closely the bargaining game in Gale (1987).

3 Defining equilibrium

We consider a Markov perfect, symmetric equilibria. In such equilibrium, the state of an individual agent at the beginning of round t is fully captured by the endowment-belief pair $(x,\delta)\in R^2\times[0,1]$, representing the agent's asset endowment $x=\left(x^1,x^2\right)$ and the probability δ he assigns to signal s_1 at the beginning of the period. An informed agent of type i begins life with the endowment-belief combination $(x_{i,0},1)$ or $(x_{i,0},0)$, i.e., his belief δ is either equal to 1 or 0 as he knows the signal s_1 and the probability of the true state. Throughout the game, the belief of an informed agent does not change. An uninformed agent of type i begins his life with $(x_{i,0},1/2)$, i.e., his initial belief is $\delta=1/2$.

The behavior of an agent at time t is described by two functions:

$$\sigma_t^p : R^2\times R_+^2\times[0,1]\rightarrow[0,1]$$

and

$$\sigma_t^r : R^2\times R_+^2\times[0,1]\rightarrow[0,1]\,.$$

The first function $\sigma_t^p(z|x,\delta)$ describes the strategy when an agent is chosen to propose an offer in a match at time t. It denotes the probability that an agent with an endowment-belief combination (x,δ) makes the offer z. The second function $\sigma_t^r(z|x,\delta)$ describes the strategy when an agent is chosen to receive an offer in a match at time t. It denotes the probability that an agent with an endowment-belief combination (x,δ) accepts the offer z. An agent strategy is then described by $\sigma=\{\sigma_t^p,\sigma_t^r\}_{t=0}^{\infty}$.

The dynamics of individual beliefs are described by two functions:

$$J_t^p : [0,1] \times R^2 \times \{0,1\} \to [0,1],$$

and

$$J_t^r : [0,1] \times R^2 \to [0,1].$$

If an agent with a belief δ is selected to propose an offer at time t, makes an offer z, the response is $r \in \{0,1\}$ ($r=0$ means that the offer is rejected, and $r=1$ means that the offer is accepted), and his belief is updated to $\delta' = J_t^p(\delta, z, r)$. If he is selected to receive the offer at time t, receives an offer z, his belief is updated to $\delta' = J_t^r(\delta, z)$.

When all players play the same strategy σ, we can construct the cross-sectional distribution of endowment-belief pairs (x, δ) across all agents, at each period t. This distribution depends on the signal s and is denoted for time t by $\Gamma_t(x, \delta | s)$. The signal s determines the initial distribution $\Gamma_0(.|s)$. For example, when the signal is s_1, the initial distribution is as follows: a mass $(1-\alpha) f_i$ of uninformed agent with the endowment-belief combination $(x_{i,0}, 1/2)$ and a mass αf_i of informed agents with the endowment-belief combination $(x_{i,0}, 1)$, for each $i \in \{1, ..., N\}$. Considering all the possible matches of agents in period 1, and the equilibrium play, described by σ_1^p and σ_1^r, we can derive the distribution of (x, δ) at the end of period 1, $\Gamma_1(.|s_1)$. Proceeding recursively, we can then derive $\Gamma_t(.|s_1)$ for all following periods.

Given the symmetry of the environment, we will focus on equilibria where strategies, updating rules, and distributions are symmetric across states s_1 and s_2. This means that the agents' behavior is identical if we interchange state s_1 for s_2, asset 1 for asset 2, and δ for $1-\delta$, all at the same time. For example, for strategy $\sigma_t^p(z|x,\delta)$ we require that $\sigma_t^p(z|x,\delta) = \sigma_t^p(z'|x',\delta')$ if $z_1 = z_2'$, $z_2 = z_1'$, $x_1 = x_2'$, $x_2 = x_1'$ and $\delta' = 1-\delta$. This symmetry across states requirement is different from the standard symmetry requirement that all agents with the same characteristics behave in the same manner which we also assume. Throughout the paper, we will use symmetry to mean symmetry across states, whenever there is no confusion. We now define an equilibrium in our model

Definition 1 *A Markov perfect Bayesian symmetric equilibrium is given by a strategy σ, a sequence of belief updating rules $\{J_t^p, J_t^r\}_t$, and a sequence of measures describing the cross-sectional distribution of beliefs and endowments $\{\Gamma_t(.|s)\}_t$ for $s \in \{s_1, s_2\}$, such that:*
(i) σ is optimal for an individual agent when his opponent at each round is randomly drawn from $\Gamma_t(.|s)$ and plays σ;
(ii) the sequence $\{\Gamma_t(.|s)\}_t$ is consistent with all agents playing the strategy σ,
(iii) the updating rules J_t^p and J_t^r are consistent with Bayes' rule whenever possible and $J_t^r(\delta, z) = J_t^p(\delta, z, r) = \delta$ if $\delta \in \{0,1\}$;
(iv) strategies, updating rules, and distributions are symmetric across states.

Relative to a standard definition of Bayesian equilibrium, we impose an additional natural restriction in the part (iii) of the above definition: informed agents' beliefs remain unchanged even after observing off-the-equilibrium path offers.

We now construct a probability space which represents uncertainty faced by an individual agent when all agents follow a given strategy σ. At time t, an agent with the endowment-belief combination (x, δ) is matched with an agent $(\tilde{x}, \tilde{\delta})$ with probability $\Gamma_t(\tilde{x}, \tilde{\delta}|s)$. Next, the agent is selected to either propose or recieve the offer with the probabilities 1/2 and 1/2. If he is chosen to propose an offer – he gives an offer z with probability $\sigma_t^p(z|x, \delta)$. His offer is accepted with the probability

$$\chi_t(z|s) = \int \sigma^r(z|\tilde{x}, \tilde{\delta}) d\Gamma_t(\tilde{x}, \tilde{\delta}|s).$$

If the agent is chosen to receive an offer at time t, he receives the offer z with probability

$$\int \sigma_t^p(z|\tilde{x}, \tilde{\delta}) d\Gamma_t(\tilde{x}, \tilde{\delta}|s)$$

and will accept with probability $\sigma^r(z|x, \delta)$.

The construction above defines a probability space $(\Omega, \mathcal{F}, P)$. The set Ω is the set of possible individual histories, i.e., the elementary event ω is given by the signal s, the initial endowment and information of the agent, and all his history of plays from period 0 to the random final date T. Let the filtration $\mathcal{F}_0 \subset \mathcal{F}_1 \subset \mathcal{F}_2 ... \subset \mathcal{F}$ describe the information sets of the agent at the beginning of each period t, immediately before nature selects whether the game ends or not. The two stochastic processes $x_t(\omega)$ and $\delta_t(\omega)$ describe the dynamics of an individual agent's endowments and beliefs in equilibrium. By construction, $x_t(\omega)$ and $\delta_t(\omega)$ are $\mathcal{F}_t$-measurable. Notice that the probability space $(\Omega, \mathcal{F}, P)$ is closely related to the cross sectional distributions $\Gamma_t(x, \delta|s)$. In particular, $\Gamma_t(x, \delta|s)$ must satisfy

$$\Gamma_t(x, \delta|s) = P(\omega : x_t(\omega) = x, \delta_t(\omega) = \delta|s).$$

4 Characterization of the equilibrium outcomes

In this section, we provide a characterization of the equilibrium in the long run, i.e., along the path when the game has not ended. We first consider the behavior of informed agents and show that they equalize their marginal rates of substitution in the long run. Then we show that the uninformed agents have the same marginal rate of substitution as the informed agents. This result implies that the uninformed agents can either construct a trade that allows them to learn the state arbitrarily well or that there are no gains from trade from the beginning of the game. This result implies that the value of information is zero in the long run, similar to the full revelation in Grossman and Stiglitz (1980). We then show that if the allocation is not Pareto efficient initially, then there are informational rents, contrasting with the classical results of Grossman and Stiglitz (1980). We finish the section with characterization of the dynamics of trade. The difficulty with analyzing this problem that we overcome is that the cross sectional distribution of beliefs is changing along the equilibrium path and is endogenous to trades that agents make or can potentially make. The agent when deciding to trade needs to know not only his belief of the state of the world, but also the beliefs of other agents, as well as forecast how they will evolve. Our proofs are constructive as we show that if the above results are not true then the agents can construct a deviating trade in which they will be able to improve their utility.

4.1 Walrasian benchmark

We briefly mention the Walrasian (centralized) benchmark to which we compare our characterization of the equilibrium. Specifically, consider a pure endowment ecomony with the same initial conditions as in our Perfect Bayesian Equilibrium, that is with I types with the initial endowments $x_{i,0} \in R_+^2$, and $\sum_{i=1}^I f_i x_{i,0}^j = 1$ for $j = 1, 2$ and a fraction α of each type observing the signal s, while fraction $1 - \alpha$ of each type having a prior $1/2$ about the realization of the signal. The rational expextation equilibrium consists of prices $p(s) \in R_+^2$ and allocations $\{x_i^*(s, \delta)\}_{i=1}^I$ where $\delta \in \{1/2, 1\}$ if $s = s_1$ and $\delta \in \{1/2, 0\}$ if $s = s_2$ s.t. $x_i^* = \arg\max_{p(s)x_i(s,\delta) \le p(s)x_{i,0}} E\{u(x)|p(s), \delta\}$ and $\sum_{i=1}^I f_i x^{j*}(s, \delta) = 1$ for $j = 1, 2$. It is a well known result that equilibrium prices fully reveal the information and all agents with the same initial endowment receive the same equilibrium allocation independently of whether they are informed or uninformed. Moreover, it is easy to show that the prices are equal to the ratio of probabilities:

$$\frac{p^1(s)}{p^2(s)} = \frac{\phi(s)}{1 - \phi(s)}.$$

Moreover, equilibrium allocations are efficient and agents equalize the allocations of goods 1 and 2: $x^{1*}(s, \delta) = x^{2*}(s, \delta)$.

4.2 Preliminary considerations

In this section we define the per-period, and lifetime utility of agents. We then show that both the lifetime utility and the beliefs converge in the long run, i.e., along the path when the game has not ended. The proof is done by applying the appropriate martingale convergence theorems.

An agent with the belief δ about the signal s_1 assigns probability $\pi(\delta)$ to the realization of state S_1, where

$$\pi(\delta) \equiv \delta\phi + (1 - \delta)(1 - \phi).$$

Note that for the informed agent, the belief δ is always equal to either 0 or 1, and $\pi(\delta)$ is then equal to either ϕ or $1 - \phi$, respectively. If the game ends, an agent with the endowment-belief pair (x, δ) receives the expected payoff

$$U(x, \delta) \equiv \pi(\delta)u(x^1) + (1 - \pi(\delta))u(x^2).$$

Using the stochastic processes x_t and δ_t, we can then define a stochastic process u_t for equilibrium expected utility of an agent if the trading game ends at the beginning of round t,

$$u_t(\omega) \equiv U(x_t(\omega), \delta_t(\omega)).$$

Let v_t denote the expected lifetime payoff of an agent at the beginning of round t,

$$v_t \equiv (1-\gamma)E\left\{\sum_{s=t}^{\infty}\gamma^{s-t}u_s \mid \mathcal{F}_t\right\}. \tag{2}$$

The next two lemmas establish that both the beliefs δ_t and the values v_t are bounded martingales and converge in the long run.

Lemma 1 *Let δ_t be the equilibrium sequence of beliefs. Then there exists a random variable δ^∞ such that*

$$\lim_{t\to\infty}\delta_t(\omega) = \delta^\infty(\omega) \ \ a.s.$$

Proof. Notice that the beliefs $\delta_t(\omega)$ are evaluated along the equilibrium path, so they must always be consistent with Bayes' rule. The law of iterated expectations implies that δ_t is a martingale,

$$\delta_t = E\left[\delta_{t+1}|\mathcal{F}_t\right].$$

Since δ_t is bounded in $[0,1]$, the result follows from the Martingale Convergence Theorem. ■

Lemma 2 *There exists a random variable $v^\infty(\omega)$ such that*

$$\lim_{t\to\infty}v_t(\omega) = v^\infty(\omega) \ \ a.s.$$

Proof. Note that an agent always has the option to reject any offers and not to make any offers from period t onwards, and wait the end of the game to consume x_t. This implies that

$$u_t \le E\left[v_{t+1} \mid \mathcal{F}_t\right]. \tag{3}$$

Equations (2) implies that

$$v_t = (1-\gamma)u_t + \gamma E\left[v_{t+1} \mid \mathcal{F}_t\right]. \tag{4}$$

Combining these results gives

$$v_t \le E\left[v_{t+1} \mid \mathcal{F}_t\right],$$

which shows that v_t is a submartingale. It is bounded above because the utility function $u(\cdot)$ is bounded above , therefore it converges by the Martingale Convergence Theorem. ■

Note that these results hold for both informed and uninformed agents.

Let us also define the function

$$\mathcal{M}(x,\delta) \equiv \frac{\pi(\delta)\,u'(x^1)}{(1-\pi(\delta))u'(x^2)},$$

which gives the *ex ante* marginal rate of substitution between the two assets for an agent with the endowment-belief pair (x,δ). We use the notation $\delta^I(s)$ to denote the belief of informed agents after signal s, so $\delta^I(s_1) = 1$ and $\delta^I(s_2) = 0$.

4.3 Long run characterization: informed agents

We now proceed to characterize the long run properties of the equilibrium outcomes. In this section we show the marginal rates of subsitution of informed agents converges a.s. to the same constant. The intuition for this result is that if it were not true, informed agents in the long run, once their values converges, could constract an offer that would be accepted by other informed agents with a different marginal rate of substitution, improving their utilities above the equilibrium payoffs. This argument is a modification of the argument of the one for the decenralized markets with full information as in Gale (1995).

Proposition 1 ***(Convergence of MRS for informed agents)*** *There exist two positive scalars $\kappa(s_1)$ and $\kappa(s_2)$ such that, conditional on each $s \in \{s_1, s_2\}$, there is a vanishing mass of informed agents with marginal rate of substitution different from $\kappa(s)$:*

$$\lim_{t\to\infty} P\left(|(\mathcal{M}(x_t,\delta_t)) - \kappa(s)| > \varepsilon, \delta_t = \delta^I(s)\,|s\right) = 0 \text{ for all } \varepsilon > 0. \tag{5}$$

Proof. We provide a sketch of the proof here and leave the complete proof for the appendix. Without loss of generality assume that the state $s = s_1$. By Lemma 2 when t is sufficiently large:

$$U(x_t(\omega), 1) \approx v_t(\omega)$$

for almost all realizations of ω. Let $\Omega^{I,1}$ be the subset of histories of informed agents in state s_1, i.e., those ω such that $\delta_t(\omega) = 1$.

If (5) is violated it is possible to construct two sets $A_t, B_t \subset \Omega^{I,1}$, both of positive measure, such that if $\omega \in A_t$ and $\tilde{\omega} \in B_t$ the difference between the marginal rates of substitution of two informed agents with $\omega \in A_t$ and $\tilde{\omega} \in B_t$ is at least $\zeta > 0$, i.e.,

$$\frac{\phi u'(x_t^1(\omega))}{(1-\phi)u'(x_t^1(\omega))} < \frac{\phi u'(x_t^1(\tilde{\omega}))}{(1-\phi)u'(x_t^1(\tilde{\omega}))} - \zeta.$$

But then agent with a history $\omega \in A_t$ can offer a small trade $z = (-\varepsilon, p\varepsilon)$ at the price

$$p = \frac{\phi u'(x_t^1(\omega))}{(1-\phi)u'(x_t^1(\omega))} + \frac{\zeta}{2},$$

i.e., he proposes to acquire ε units of the first asset by giving $p\varepsilon$ units of the second asset.

The utility of the agent ω is higher if his offer is accepted since

$$\begin{aligned} U(x_t^1(\omega) - z, \phi) &= \phi u(x_t^1(\omega) + \varepsilon) + (1-\phi)u(x_t^2(\omega) - p\varepsilon) \\ &\approx U(x_t^1(\omega), \phi) + \left[\phi u'(x_t^1(\omega)) - (1-\phi)pu'(x_t^2(\omega))\right]\varepsilon \\ &\approx v_t + (1-\phi)pu'(x_t^2(\omega))\frac{\zeta}{2}\varepsilon. \end{aligned}$$

By choosing t sufficiently large and ε sufficiently small, we can make the approximation errors in the above equation sufficiently small, so that such trade improves the utility of the first type of agents, $U(x_t^1(\omega) - z, \phi) > v_t$. All the formal steps are in the appendix.

The agent with the event $\tilde{\omega} \in B_t$ is also better off by an similar argument. Therefore, all informed agents in B would accepted the offer. Since there is a positive probability for agent ω to meet an agent $\tilde{\omega} \in B$, his utility from following this strategy is strictly higher than from following the equilibrium strategy. Therefore, we obtain a contradiction.

This argument shows that there is a sequence of $\kappa_t(s_1)$, possibly varying over time, to which marginal rate of substitution of the informed agents converges. Since in any equilibrium each agent faces a positive probability that no trade is make (e.g. if he meets an agent of the same type (x, δ)), $\kappa_t(s_1)$ is constant over time, $\kappa_t(s_1) = \kappa(s_1)$, completing our proof. ■

There are two important points to note in the argument above. First, there may be a mass of uninformed agents who also potentially accept the offer of z, but this only increases the probability of acceptance, which further improves the utility of the agent ω. Second, the offer of z is not necessarily the best offer and the the proof described not necessarily the best strategy the agent ω can follow. Potentially, there are even better sequence of offers that the agent ω can make to improve his utility. However, considering the offer z is sufficient to show that the agent can get higher utility that v_t and arrive to a contradiction. This second remark is important as it is the strategy of proof in many of the results that follow.

We can show the following corollary that ensures convergence of allocations by informed agents.

Corollary 1 *For all informed agents, the process $\{x_t\}$ almost surely converges to a constant.*

Proof. Suppose this is not the case. From Lemma 1 there must exist two subsequences of $x_t(\omega)$, one converging to to x' and the other converging to x'', both leading to the same marginal rates of substitution:

$$\frac{\phi u'(x^{1\prime})}{(1-\phi)u'(x^{2\prime})} = \frac{\phi u'(x^{1\prime\prime})}{(1-\phi)u'(x^{2\prime\prime})}.$$

This is possible only if $x' > x''$ or $x' < x''$ which, however, violates Lemma 2. ■

4.4 Long run characterization: uninformed agents

We now turn to characterizing equilibrium for the uninformed agents. The main difficulty in the analysis is that the uninformed agents, upon receiving offers or upon having their offers accepted or rejected, might change their beliefs. Thus agents, who might be willing to accept some offer ex-ante before updating their beliefs, might reject it after an update. An additional complexity comes from the fact that such beliefs for an arbitrary offer are usually not pinned down by Bayes rule, since such offers may not occur in equilibrium. For these reasons we need to pursue proof strategies that differ from those that we used to characterize behavior of the informed agents.

Our arguments are based on the finding strategies that allow uninformed to learn the signal s at an arbitrarily small cost in the long run. The existence of such strategies implies that either agents indeed eventually learn the signal, or the benefits of such learning is zero. In the next section we'll show that both of these cases imply that equilibrium allocations converge to the efficient ones in the long run.

The particular learning strategies for the uninformed we consider depend on wether the marginal rates of substitution of informed agents converge to the same value in both states of the world, $\kappa(s_1) = \kappa(s_2)$, or to different values. For this reason we split this section in two parts, one that analyses case $\kappa(s_1) \neq \kappa(s_2)$, and the other one for the case $\kappa(s_1) = \kappa(s_2)$. We'll also show that the abilty of the uninformed to learn the signal s at a small cost when $\kappa(s_1) = \kappa(s_2)$ implies that this case cannot occur in equilibrium.

4.4.1 The case $\kappa(s_1) \neq \kappa(s_2)$

We begin by considering the case in which the marginal rates of substitutions for informed agents converge to different numbers in states s_1 and s_2: $\kappa(s_1) \neq \kappa(s_2)$.

Proposition 2 ***(Convergence of MRS for uninformed agents)*** *Suppose $\kappa(s_1) \neq \kappa(s_2)$. Let x_t be an equilibrium process for any (informed or uniformed) agent. For each $s \in \{s_1, s_2\}$,*

$$\lim_{t\to\infty} P\left(\left|\frac{\phi(s)u'(x_t^1)}{(1-\phi(s))\,u'(x_t^2)} - \kappa(s)\right| \mid s\right) = 0. \tag{6}$$

Note that the Proposition states that *if* one is to evaluate marginal rate of substitution of any agent, that MRS is equal to $\kappa(s)$. The actual marginal rates of substitution of uninformed is $\pi(\delta_t)u'(x_t^1)/(1-\pi(\delta_t))u'(x_t^2)$ and it might be different from $\kappa(s)$ if δ does not converge to 1 if $s = s_1$ or 0 if $s = s_2$.

Proof. We provide a sketch of the proof here and leave the complete proof for the appendix. In this sketch we'll use all the arguments for the steady state assuming that all the agents have converged to their long run values, while in the complete proof shows that the same arguments are true for t sufficiently large.

Since symmetry implies $\kappa(s_1) = 1/\kappa(s_2)$, assume without loss of generality that $\kappa(s_1) > 1 > \kappa(s_2)$. Consider an offer $z = (\varepsilon, -\varepsilon)$.

Observation 1. Offer z is accepting by all informed agents if $s = s_1$. Since by assumption all agents converge to their long run values, for any informed agent whose allocation converges to $\tilde{x}$,

$$\begin{aligned} U(\tilde{x}+z, \phi(s_1)) &\approx U(\tilde{x}, \phi(s_1)) + (\kappa(s_1)-1)(1-\phi(s_1))u'(\tilde{x}^2)\varepsilon \\ &> U(\tilde{x}, \phi(s_1)) = \tilde{v} \end{aligned}$$

where $\tilde{v}$ is the long run value to this payoff of the informed converged.

Observation 2. Offer z is not accepted by any informed agent if $s = s_2$. If this were the case, and some fraction of informed agent accepted offer z in state $s = s_2$, symmetry of equilibrium would imply that the same fraction would accept an offer $-z$ in state $s = s_1$. But then from Observation 1 in the long run both offers z and $-z$ would strictly improve utility of informed agents, which would contradict the assumption that all informed converge to their long run values and cannot improve their utility further.

Observation 3. Offer z is not accepted by any uninformed agent with $\delta^{\infty} \in (0,1)$ in any state.[3] Suppose otherwise and a positive fraction of uninformed accept z in some state s. Since $\delta^{\infty} \in (0,1)$ for such agents, they must have observed histories that have positive probability of occuring in both states, and therefore there must be a positive fraction of uninformed in both states that accept z in the long run in both states. Symmtry argument analogous to the one made in observation 2 implies that informed agents in states s_2 are strictly better off at making an offer $-z$. Such offer is accepted by the uninformed and allows informed to strictly improve their utility, which is inconsistent with the assumption that all equilirbium payoffs converged.

Suppose that the proposition is not true, so that without loss of generality (6) is not satified for a positive mass of uninformed agents in state $s = s_1$. For concreteness, let $\frac{\phi(s_1)u'(x^1)}{(1-\phi(s_1))u'(x^2)} < \kappa(s_1)$, where $x = \left(x^1, x^2\right)$ is the value to which some subsequence of x_t of such uninformed converges. Since Lemma 1 shows that $\delta_t \to \delta^{\infty}$, Proposition 1 implies that $\delta^{\infty} \in (0,1)$. Let v^{∞} be the long run value to which the payoff of such uninformed converges. Note that $v^{\infty} = U(x, \delta^{\infty})$.

In what follows we construct a strategy for uninformed with asset position x. Suppose that upon reaching x, the uninformed makes the following deviation. If he can make an offer, he proposes an offer $z_{\varepsilon} = (\varepsilon. -\varepsilon)$. If he cannot, or if his offer is rejected, he rejects all offers after that and makes $(0,0)$ offer whenever he is a proposer, so that his position in those contigencies remains x. If his offer is accepted, in the next round he makes an offer $z_{\theta} = (\theta, -\zeta\theta)$ where $\zeta = \frac{1}{2}\left(\kappa(s) + \frac{\phi(s_1)u'(x^1)}{(1-\phi(s_1))u'(x^2)}\right)$ and $\theta > 0$ and small. In choosing θ and ε we set ε to be small relative to θ, i.e. $\varepsilon = o(\theta)$. He does not make trades in subsequent rounds, or if he is not a proposer.

With this strategy, if $s = s_1$, the uninformed agent in two round has an allocation $x - z_{\varepsilon} - z_{\theta}$ with probability $\frac{1}{4}\alpha^2\gamma$, which happens in the case that he is chosen to make offers in two subsequent rounds, and in both cases he is matched with informed agents. For all other contigencies his allocation is either x or $x - z_{\varepsilon}$, which for any δ implies $U(x - z_{\varepsilon}, \delta) = U(x, \delta) + o(\theta)$ since $\varepsilon = o(\theta)$. This implies that the ex-ante payoff of the uninformed agent with allocation x from following this strategy is

$$\begin{aligned} & U(x, \delta^{\infty}) + \left(\kappa(s_1) - \frac{\phi(s_1)u'(x_t^1)}{(1-\phi(s_1))\,u'(x_t^2)}\right)(1-\phi(s_1))\,u'(x_t^2)\theta + o(\theta) \\ > \quad & U(x, \delta^{\infty}) = v^{\infty} \end{aligned}$$

[3] Note that the definition of equilibrium implies that any uninformed for whom $\delta^{\infty} \in \{0,1\}$ behaves like an informed in those states.

One can see it is true by the following:

$$
\begin{aligned}
& \frac{\alpha}{2}\delta^{\infty}\left(\gamma\frac{\alpha}{2}U(x-z_{\varepsilon}-z_{\theta},1)+\gamma\left(1-\frac{\alpha}{2}\right)U(x-z_{\varepsilon},1)+(1-\gamma)U(x-z_{\varepsilon},1)\right) \\
& +(1-\frac{\alpha}{2})\delta^{\infty}U(x,1)+(1-\delta^{\infty})U(x,0) \\
= \quad & \frac{\alpha}{2}\delta^{\infty}\left(\gamma\frac{\alpha}{2}U(x-z_{\theta},1)+\gamma\left(1-\frac{\alpha}{2}\right)U(x,1)+(1-\gamma)U(x,1)\right) \\
& +(1-\frac{\alpha}{2})\delta^{\infty}U(x,1)+(1-\delta^{\infty})U(x,0)+o(\theta) \\
= \quad & \delta^{\infty}U(x,1)+(1-\delta^{\infty})U(x,0)+\frac{1}{4}\alpha^{2}\gamma\delta^{\infty}\left(\kappa(s_1)-\frac{\phi(s_1)u'(x_t^1)}{(1-\phi(s_1))\,u'(x_t^2)}\right)(1-\phi(s_1))\,u'(x_t^2)\theta+o(\theta) \\
> \quad & U(x,\delta^{\infty})=v^{\infty}
\end{aligned}
$$

This shows that an uninformed can further improve his utility, leading to a contradiction to the assumption that his long run value is reached. ■

The most subtle part of the proof is the arguments behind Observation 3. Once we established that for any small ε an offer $(\varepsilon, -\varepsilon)$ is accepted by only informed agents in state s_1 and by no agents in state s_2, it is easy to describe how agent can improve his utility. By making an offer z can learns that the signal is s_1 with probability $\delta\alpha$, and then can improve his utility by trading with informed agents and equilizing his marginal rate of substitution similar to the argument in the proof of Proposition 1. Since ε can be chosen arbitrarily small, the utility loss for all other realizations can be made less, than the expected gain from learning and trading with the informed, which leads to the contradition the possibility that uninformed's marginal rate of subsitution can be different from $\kappa(s)$ if evaluated at the objective probabilities.

4.4.2 The case $\kappa(s_1) = \kappa(s_2)$

When $\kappa(s_1) = \kappa(s_2)$, learning the signal is more difficult for the uninformed than in the proofs in the previous section. Informed agents in both states have the same marginal rates of substitutions, and their strategies for the small offers might be the same. Consider, for example, any uninformed agent whose marginal rate of subsitution $\mathcal{M}(x_t, \delta_t)$ does not converge to $\kappa(s)$, so that, for example, there is a subsequence of $\mathcal{M}(x_t, \delta_t)$ that converges to $\mathcal{M} < \kappa(s)$. Suppose once uninformed agent's marginal rate of substitution is close to $\mathcal{M}$, such an agent deviates from his equilibrium strategy and makes offers $z = (\varepsilon, -\varepsilon\zeta)$ where $\zeta = \frac{1}{2}(\kappa(s) + \mathcal{M})$ when he is a proposer. If the probabilities of acceptance of offer z aresufficiently different in the two states, the agent learns s. Once he knows s, the agent can further improve his utility by following strategies described in the proof of Proposition 2. On the other hand, if these probabilities are sufficiently similar, his updated beliefs should be close to his beliefs δ_t, but then the trade z is constructed in such a way that it increases his expected utility if it is evaluated at the subjective beliefs δ_t. Both of these cases lead to a contradiction that agent's long run payoff converged. We formally state the proposition below and leave the proof to the appendix.

Proposition 3 *Suppose $\kappa(s_1) = \kappa(s_2) = 1$ and suppose that for some $s \in \{s_1, s_2\}$, $\varepsilon > 0$ and $\zeta > 0$,* *there is an infinite sequence of periods* $\{t_k\}_{k=1}^{\infty}$ *such that*

$$P\left(\left|\mathcal{M}\left(x_{t_k},\delta_{t_k}\right)-\kappa(s)\right|>\zeta\mid s\right)>\varepsilon$$

for $k=1,2,\ldots$. *Then*

$$\lim_{k\to\infty}\left(t_{k+1}-t_k\right)=\infty. \tag{7}$$

There are two differences of Proposition 3 with Proposition 2. First, here we show convergence of agent's marginal rates of substitution evaluated at his beliefs, $\pi(\delta_t)u'(x_t^1)/(1-\pi(\delta_t))u'(x_t^2)$, rather than at actual probabilities $\phi(s)u'(x_t^1)/(1-\phi(s))u'(x_t^2)$. Second, here we show convergence in a weaker sense than convergence in probability in the previous section. Namely, we show that marginal rates of substitution converge for all periods except for a sequence of periods $\{t_k\}_{k=1}^{\infty}$, with $t_{k+1}-t_k\to\infty$. That is, there may be periods where marginal rates of substitution differ across agents, but these periods become increasingly rare as time progresses. It turns out that this result is strong enough to show that there are no equilibria with $\kappa\left(s_1\right)=\kappa\left(s_2\right)$.

Proposition 4 *There is no symmetric equilibrium with* $\kappa\left(s_1\right)=\kappa\left(s_2\right)$.

Proof. Once again, we sketch the arguments, leaving technical details for the appendix. Symmetry implies that $\kappa\left(s_1\right)=\kappa\left(s_2\right)=1$ and Proposition 3 that for t large marginal rates of substitution of almost all agents are close to 1. Feasibility implies that

$$\int x_t^1 d\Gamma(\omega|s)=\int x_t^2 d\Gamma(\omega|s)=1 \tag{8}$$

Suppose $s=s_1$ and consider any agent with $\delta_t(\hat{\omega})\in(0,1/2]$. Interiority of beliefs implies that there is positive probability of observing the same history in both states. Symmetry of equilibrium and Bayes rule implies that in equilibrium for any $\hat{\omega}$ s.t. $\delta_t(\hat{\omega})\in(0,1)$ there is $\tilde{\omega}$ s.t. $\delta_t(\tilde{\omega})=1-\delta_t(\hat{\omega})$ and that there are two agents whose asset positions are symmetric to each others in all dates: $\left\{\left(x_i^1(\hat{\omega}),x_i^2(\hat{\omega})\right)\right\}_{i=1}^{t}=\left\{\left(x_i^2(\tilde{\omega}),x_i^1(\tilde{\omega})\right)\right\}_{i=1}^{t}$. Moreover if $\delta_t(\hat{\omega})\leq 1/2$, than then measure of agents with history $\tilde{\omega}$ is strictly greater than the measure of agents with a history $\hat{\omega}$.[4] Then for any measure of agents with history $\hat{\omega}$ we can find the same measure of agents with a history $\tilde{\omega}$. We can form a new distribution $\tilde{\Gamma}\left(\omega|s_1\right)$ by setting $\tilde{\Gamma}(\omega|s_1)=\Gamma(\omega|s_1)-\Gamma(\hat{\omega}|s_1)$ if $\omega\in\{\hat{\omega},\tilde{\omega}\}$ and $\tilde{\Gamma}(\omega|s_1)=\Gamma(\omega|s_1)$ otherwise. By construction $\int x_t^1 d\tilde{\Gamma}(\omega|s_1)=\int x_t^2 d\tilde{\Gamma}(\omega|s_1)$. Repeating this procedure for all ω with $\delta_t(\omega)\in(0,1/2]$, we construct a distribution $\Gamma^*(\omega|s_1)$ s.t.

$$\int x_t^1 d\Gamma^*(\omega|s_1)=\int x_t^2 d\Gamma^*(\omega|s_1) \tag{9}$$

and $\Gamma^*(\omega|s_1)=0$ for all ω s.t. $\delta_t(\omega)=0$. Since for almost all ω s.t. $\Gamma^*(\omega|s_1)>0$, $\pi(\delta_t)u'(x_t^1)/(1-\pi(\delta_t))u'(x_t^2)$ converges to 1, and $\pi(\delta_t)/(1-\pi(\delta_t))>1$, it must be true that $x_t^1(\omega)>x_t^2(\omega)$ which contradicts (9). ■

[4]See the proof of Lemma 11 in the appendix for the formal proof of these properties.

4.5 Efficiency and informational rents

The characterization of the behavior of informed and uninformed agents in the previous section allows us to derive the key results about the long run efficiency and the value of information.

Theorem 1 ***(Efficiency in the long run)*** *Equilibrium allocations converge to efficient allocations in the long run, i.e.,*

$$\lim_{t\to\infty} P\left(\left|x_t^1 - x_t^2\right| > \varepsilon\right) = 0 \text{ for all } \varepsilon > 0. \tag{10}$$

For any $s \in \{s_1, s_2\}$ the long run marginal rate of substitution $\kappa(s)$ is equal to the ratio of the conditional probabilities of states S_1 and S_2,

$$\kappa(s) = \phi(s)/(1-\phi(s)).$$

Proof. Proposition 4 rules out the case $\kappa(s_1) = \kappa(s_2)$. First suppose that $\kappa(s) > \phi(s)/(1-\phi(s))$ for some s. Then Propositions 1 and 2 imply that $\lim_{t\to\infty}\left(x_t^1 - x_t^2\right) < 0$ a.s. This, however, violates (8). We rule out the case $\kappa(s) < \phi(s)/(1-\phi(s))$ analogously. Since $\kappa(s) = \phi(s)/(1-\phi(s))$, 1 and 2 imply (10). ■

This theorem establishes that in the long run equilibrium allocations coincide with a rational expectation equilibrium defined in Section 4.1 for *some* initial allocations. It does not show whether starting with the *same* initial allocations centralize rational expectation equilibrium and the decentrilized matching environment we consider converge to the *same* long run outcomes. To futher explore whether informed agents can achieve higher payoff than uninformed agents, we define a concept *value of information*. Consider any agent in period t with a history $h^t \in \mathcal{F}_t$. Let $\delta_t(h^t)$ be beliefs of such an agent in period t, and $v(x_t|h_t)$ his expected payoff. Let $v(x_t, t|s)$ be a payoff of an agent in period t if he had endowment x_t and believed that the signal was s w.p. 1. Thus, $v(x_t, t|s)$ is a payoff of a hypothetical informed agent in period t with endowment x_t.[5] If an agent with a history h^t could costlessly learn signal s, his expected utility would increase by $\mathcal{I}(h^t) = \delta\left(h^t\right) v(x_t, t|s_1) + (1-\delta\left(h^t\right))v(x_t, t|s_2) - v(x_t|h_t)$. We call $\mathcal{I}_t$ to be the *value of information*. Since upon costlessly learning s an agent can continue to persue his equilirbium strategy, the value of information is always nonnegative, $\mathcal{I}_t \geq 0$.

The first result that follows shows that a famous no trade theorem due to Milgrom and Stokey (1982) holds in our settings, and if agents begin with Pareto efficient allocations the value of information is zero.

Theorem 2 *Suppose $x_{0,i}^1 = x_{0,i}^2$ for all i. Then there is no trade in equilibrium and $\mathcal{I}_t = 0$ for all t.*

Proof. It is a straightforward adaptation of the proof of Theorem 1 in Milgrom and Stokey (1982) once one notices that allocations are Pareto Efficient if and only if $x_i^1 = x_i^2$ for all i. ■

One of the implications of this result is that informed and uninformed agents with the same initial endowments receive the same payoff if the initial allocation is efficient. Combining the insight of Theorem 2 with the result from Theorem 1, it is easy to obtain the following corrollary

[5] In equilibrium there might not exist an informed agent with endowment x_t in period t. However, we can formally extend our game by addition a measure 0 of such agents in period t. Since they are of measure zero, such an extension does not change the equilibrium of the game, but the payoffs $v(x_t, t|s)$ become well defined.

Corollary 2 *The value of information is zero in the long run, i.e. for all s*

$$\lim_{t\to\infty} P(\mathcal{I}_t > \varepsilon|s) = 0 \text{ for any } \varepsilon > 0.$$

The question remains whether informed agents can get a higher utility than uninformed if the initial allocation is not Pareto Efficient. The next theorem shows that this is indeed the case, and the the following two sections we explore which strategies informed agent might persue that increases their utility.

Theorem 3 *Suppose $x_{0,i}^1 \neq x_{0,i}^2$ for some i. Then the value of information is positive in period 0, i.e. there exists some $\varepsilon > 0$ s.t.*

$$P(\mathcal{I}_0 \geq \varepsilon|s) > 0$$

for all s.

Proof. In the appendix. ■

5 A static example

In this section we give more insights to the behavior of agents in the model on a static example. This section useful to also provide intuition for the numerical section in which we solve a fully dynamic game.

Suppose there are two types of agents with endowments (e_h, e_l) and (e_l, e_h) of the two goods with $e_h > e_l > 0$. The environment is static. There is only one round of random matching, after which the game ends, the state of the world is realized, and each agent consumes his endowment. We further simplify the environment by assuming that the fraction α of informed agents is negligible – this implies that their presence does not affect strategies of uninformed agents.

Figure 4 and Figure 5 represents strategies of informed and uninformed players using an Edgeworth box. Note that since the only trade in equilibrium would occur between agents of different types, we can focus without loss of generality only on the situations when an agent who makes an offer has an endowment (e_h, e_l) and a receiver has an endowment (e_l, e_h).

The equilibrium behavior of uninformed agents depends on the beliefs that an uninformed agent who recieves the offer forms with respect to various out of equilibrium offers. In general, many out of equilibrium beliefs are possible which in turn implies that there are multiple equilibria in this static game. For the purposes of studying this static example, we put more structure on the out of equilibrim beliefs. Speicifically, we focus on a particular equilibrium selection, namely, on an equilibrium that satisfies the intuitive criterion of Cho and Kreps (1988). However, it is important to note that our results in the previous sections apply to all equilibiria and do not rely on any refinements or equilibrium selection.

First, consider an uniformed agent who makes an offer. This offer must maximize his utility and also ensure that an agent who receives the offer accepts such trade, i.e., it should give a weakly higher utility than the initial endowment. Since we assumed in this static example that the fraction of informed agents is negligible, the agent who receives the offer does not update his beliefs. Both of the uninformed agents value each of the assets equally, the slopes of their indifference curves on the 45 degree line are equal to 1/2 (the red dotted line is the indifference curve for the uninformed agent who receives an offer and the blue dotted line is an indifference curve for the agent who makes an offer). The red dot on Figure 4 represents such an equilibrium point.

Next, consider a strategy of the informed agent who makes an offer. Such agent received a signal that that state s_1 is more likely. Since his initial endowment of good 1 is greater than endowment of good 2, we call him a rich informed agent. His indifference curve is steeper (blue dashed line) that those of uninformed as he knows that good 1 is more valuable. The next step is to consider whether such an informed agent has incentives to make an offer different from those that the uninformed agents make, i.e., to choose a point different from the red dot in Figure 4. To proceed, as discussed above, we need to restrict the out of equilibrium beliefs of the uninformed agent who receives an offer. We focus on a commonly used restriction on the out of equilibrium beliefs due to Cho and Kreps (1988). In the context of our static example, the intuitive criterion implies the following . Consider an uninformed agent who receives an out of equilibrium offer z. If such offer gives a higher utility only to one type of agents than any other equilibrium offers, the agent who receives the offer changes his beliefs to assume that the offer is being made by that particular type.

The shaded area on the figure shows all offers that give the rich informed agent who makes it higher utility than the utility which the uninformed agent who makes it (and a higher utility than that of the poor informed agent, as we show below). For any offer that is being made in that region, the uninformed agent who recieves the offer must therefore infer that the the agent who makes the offer is a rich informed agent, and therefore that the informed agent observed signal s_1. Before making a decision whether to accept or reject such an offer, the agent who receives the offer updates his beliefs. His indifference curve then becomes steeper and is represented on the Figure 4 by a red dashed line. Therefore, the rich informed agent makes an offer that maximizes his utility in the shaded region (i.e., in the region that changes the beliefs of the counterpart) and leaves the receiver weakly better off than rejecting the offer. This is given by the blue dot on the graph. Note, that the price that an agent who receives the offer pays for good 1 (defined as a quantity of good 2 paid for a unit of good 1) is higher than: (a) the price at which uninformed agents offer good 1, and (b) than the price that informed agents would offer in an equilibrium where all agents have full information (green dot on the graph). Why is the green dot not an equilibrium? The reason is that at that point both the agents who are informed and uninformed can make the the offer – hence, the agent who recieves such offer does not change his beliefs and does not shift the indifference curve, remains on the dotted red indifference curve, and rejects the offer.

This example shows one of the possibilities for the informed agents to receive higher payoff than uninformed agents. By offering a small amount of good 1 at a high price, the informed agent can credibly signal that they have information. The reason is that only if an agent is informed that state s_1 is more likely he is willing to retain so much of good 1. Such signalling leads to the ex-post inefficiency. One can see that inefficiency by observing that the equilibrium point (the blue dot) does not lie on the

45 degree line. In contrast, the full information equilibrium (the green dot) is efficient and lies on the 45 degree line.

Figure 5 considers a case when an informed agent receives information that state 2 is more likely. If he makes any offer that signals his type (an offer in the shaded area), the uninformed agent updates his beliefs, changes the indifference curve to the one represented by the dashed line, and rejects the offer. In other words, such an offer reveals to the uninformed agent that his reservation value is higher than an ex-ante reservation value, evaluated with equal probabilities for two states of the world (i.e., the one on the dotted red indifference curve). Therefore, such offer would not be accepted by the receiver with the updated beliefs. For this reason the poor informed agent would rather replicate the strategies of the uninformed agent than reveal his type, i.e., make offers at the red dot.

This example illustrate several important features of agents' strategies which we show will remain true in a computed equilibrium of the dynamic game. First, since rich informed agents need to signal their information, it usually takes longer to reach efficient outcomes than with full information. Second, rich informed types generally prefer to sell little of their endowment early in the game, and they decrease their position slowly. Finally, poor informed sell their endowment relatively fast, taking advantage of the fact that uninformed agents do not know which asset is more valuable.

6 Numerical illustration

In this section we illustrate quantitatively the theoretical results of the paper derived so far. The analysis of this section allows to show some interesting properties of the equilibria to complement our theoretical derivations. We also contrast our results with the case when all the information is public.

Consider a per period utility of an agent to be given by

$$u(x) = -\exp(-x).$$

One can easily show that with such utility function, what matters for the agent is the difference in endowments $x^1 - x^2$. We can then collapse the state of the agent to be (x, δ, t).

Suppose there are two types of agents: half of agents starts with the endowment $(2, 0)$ and the other half starts with the initial endowment $(0, 2)$. Consider the state of the world s_1, i.e., the first good is more valuable. We call the first type of agents "rich" and the second type of agents "poor".

Figure 1 is the most interesting in this section and describes the strategies of various agents in the economy.

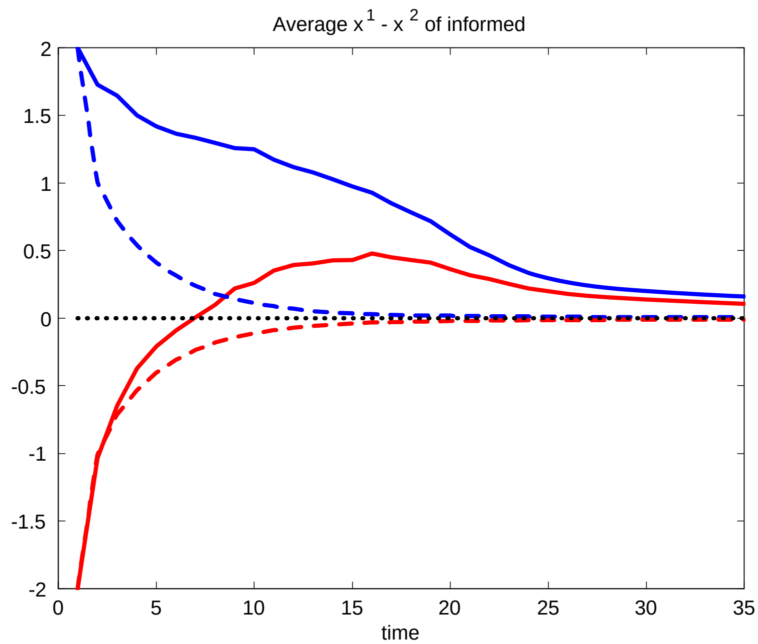

Trades of informed agents under public and private information

The dotted lines in the picture show the case of full information. The dashed red line describes the average endowment position $(x^1 - x^2)$ of the "poor" agents. We see that as such agents meet other agents (rich agents in that case), the poor agent trades to acquire the second asset and finally converges to the efficient allocation in which $x^1 = x^2$. A similar strategy is for the rich agent who sells some of his first asset, acquires the second asset and converges to the efficient allocation. The question may arise why the convergence does not occur in one period. The reason is that some agents may meet with the agents of the same type, for example if in the first period, two "rich" agents meet, there will be no trade.

A different picture arises in the setting with private information. Consider first the "poor" informed agents, i.e., those who started with the endowment $(0, 2)$. Their trades are depicted by the solid red line. As was discussed in the theoretical part, such agents know that in the long run the uninformed agents will learn the true state of the world and the terms of trade will turn against the "poor" agents. Therefore, the "poor" informed agents buy the first asset as fast as possible, and the red line rises steeply. Moreover, they buy more of the first (more valuable) assets than they will eventually end up with and "overshoot" to have $(x^1 - x^2) > 0$, and then eventually decrease their holdings to the efficient allocation. The incentives of the "rich" agents are different (solid blue line). They want to hold on to their endowment of the valuable good until the information will be revealed. These strategies also reveal how informed agents receive a positive lifetime utility from having private information.

We also see that the it takes longer to achieve efficiency in the economy with private information. For example, in period $t = 15$, in the case of public information the efficiency is essentially achieved, while for our private information economy, the agents are still far from the efficient allocations.

The next graph, Figure 2 illustrates how efficiency is achieved more slowly under the case of private information (solid line) than in the case of public information (dotted line). Here we plot the standard deviation of the difference in endowments in the economy.

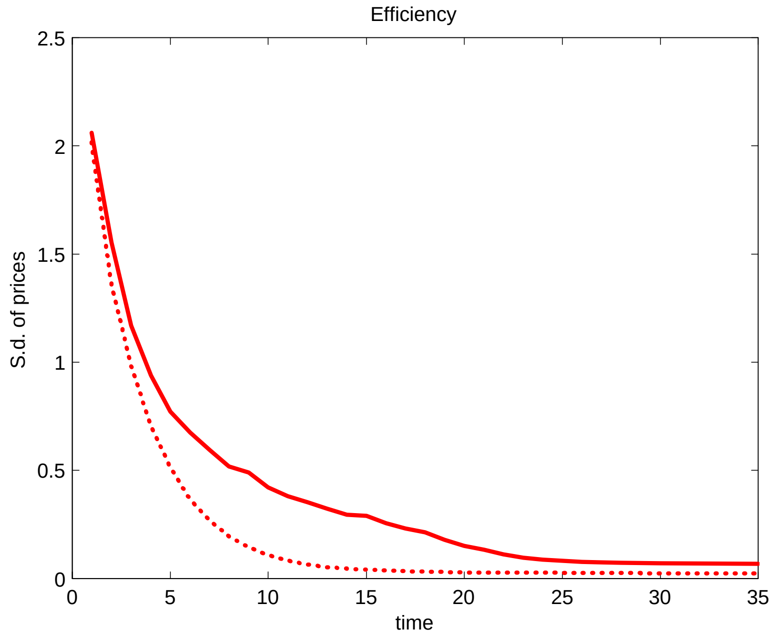

Efficiency under private and public information

Finally, we provide the graph of the volume of trades in our environment with private information, Figure 3. One could have conjectured that the volume of trade will be higher in this environment as uninformed agents strategically experiment to learn the true state of the world. This turns out not to be the case, and the volume of trade in the environment with private information (in the graph) is virtually identical to the volume in the environment with public information. The intuition for this is as follows. 'As can be seen in Figure 1, the informed rich agents wait for some time to trade, and if they trade in the short run they trade small amounts. They act as counterparties for the uninformed agents who strategically experiment with relatively small trades.

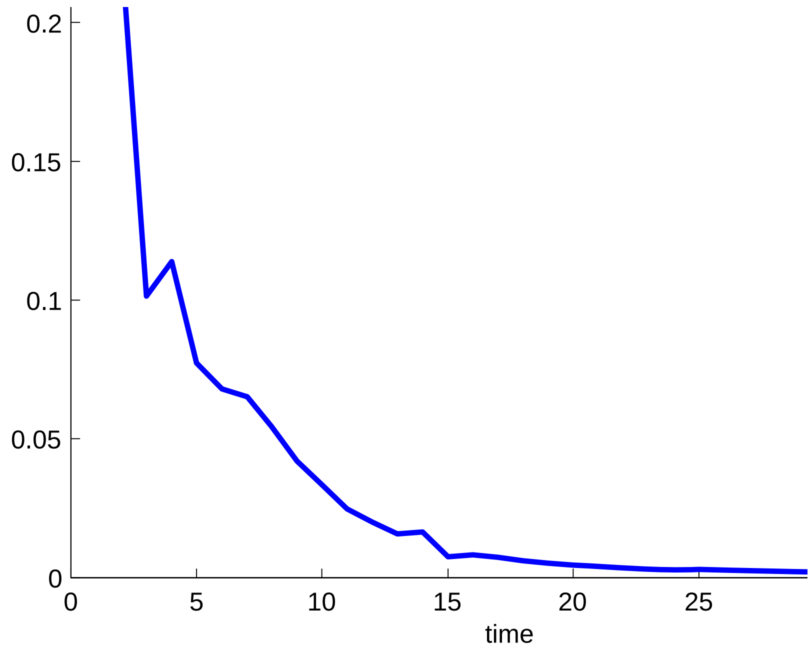

Volume of trades

7 Conclusion

We provide a theory of asset pricing in an environments which are characterized by two frictions: (1) private information; (2) decentralized trade. These frictions often go hand in hand – it is reasonable to think that the decentralized markets (such as, for example, over the counter markets) also have large amounts of private information. While the analysis of asset pricing under asymmetric information in centralized markets is well developed (see, e.g, a comprehensive examination in Brunnermeier, 2001), ours is one of the first papers to develop such theory in the decentralized environments.

Our results on convergence to the efficient allocation, learning by uninformed agents, and dynamics of trade are very general and do not rely on specific functional forms. The reason for this generality is that we employ a novel argument by constructing a trade in which the uninformed agents can experiment and learn the true state. We also provide a numerical simulation that illustrates and extends the theoretical part of the paper.

8 Appendix

(incomplete)

8.1 Technical Preliminaries

Here, we prove a number of preliminary results that will be useful throughout the appendix. Lemma 3 shows that the per period and lifetime utilities converge. Lemma 4 shows that endowments converge to a compact interior set with the probability arbitrarily close to one. This set is needed mainly for technical reasons, to ensure that maximization problems used in the proofs are well defined. Lemma 5 is a technical result from probability theory that will be useful when we want to combine a statement which holds with positive probability in the long run with a statement which holds with probability arbitrarily close to one.

The lemma that follows is a simple result that shows that the per-period utility converges to the lifetime utility for sufficiently large t both unconditionally and conditional on a given realization of the signal s. The proof uses Lemma 2 and simple algebraic manipulations.

Lemma 3 *For any* $\varepsilon > 0$

$$\lim_{t\to\infty} P\left(|u_t - v_t| \geq \varepsilon\right) = 0$$

and

$$\lim_{t\to\infty} P\left(|u_t - v_t| \geq \varepsilon | s\right) = 0$$

for $s \in \{s_1, s_2\}$.

Proof. First we show that for any $\varepsilon > 0$

$$\lim_{t\to\infty} P(|v_t - E[v_{t+1} \mid \mathcal{F}_t]| \geq \varepsilon) = 0. \tag{11}$$

As argued in Lemma 2, v_t is a bounded supermartingale and converges almost surely to v^∞. Let $y_t \equiv E[v_{t+1} \mid \mathcal{F}_t]$. Since a bounded martingale is uniformly integrable (see Williams 2005), we get $y_t - v_t \to 0$ almost surely. Then note that almost sure convergence implies convergence in probability, so $\lim_{t\to\infty} P\left(|y_t - v_t| \geq \varepsilon\right) = 0$ for all $\varepsilon > 0$. Rewrite equation (4) as

$$(1-\gamma)\, u_t = \gamma\left(v_t - E\left[v_{t+1} \mid \mathcal{F}_t\right]\right) + (1-\gamma)\, v_t.$$

This gives

$$u_t - v_t = \frac{\gamma}{1-\gamma}\left(v_t - E\left[v_{t+1} \mid \mathcal{F}_t\right]\right),$$

which combined with (11) shows that $\lim_{t\to\infty} P\left(|v_t - u_t| \geq \varepsilon\right) = 0$.

To prove the second part of the lemma notice that

$$P\left(|v_t - u_t| \geq \varepsilon\right) = \sum_{k=1,2} P\left(s_k\right) P\left(|v_t - u_t| \geq \varepsilon | s_k\right),$$

where $P(s_k) = 1/2$ for $k = 1, 2$. Therefore, given any $\eta > 0$, $P(|v_t - u_t| \geq \varepsilon) < \eta$ implies $P(|v_t - u_t| \geq \varepsilon | s_k) < 2\eta$ for $k = 1, 2$. ■

The next lemma uses market clearing and the convergence of the lifetime utility v_t to show that we can always find a compact set X in the interior of R^2_+, such that the allocations x_t are in that set with a sufficiently high probability in the long run. In the rest of the appendix, we construct compact sets with this property on various occasions by using the next lemma, to set bounds on the utility gains that agents of different types derive from trades. The proof is a somewhat tedious constructive argument that uses the convergence of utilities and the market clearing conditions.

Lemma 4 *For any $\varepsilon > 0$ and for any $s \in \{s_1, s_2\}$, there exists a time T and a compact set $X \subset R^2_+$ which lies in the interior of R^2_+, such that $P(x_t \in X | s) \geq 1 - \varepsilon$ for all $t \geq T$.*

Proof. Pick any $\eta > 0$ and choose T so that $P(|u_t - v_t| > \eta | s) \leq \varepsilon/4$ for all $t \geq T$, i.e., so that u_t is sufficiently close to v_t. This can be done by Lemma 3.

Next, we use goods market clearing to show that $P(x_t^j > 8/\varepsilon | s) \leq \varepsilon/4$ for each good $j = 1, 2$. To prove this, notice that

$$1 = \int x_t^j(\omega)\, dP(\omega|s) \geq \int_{x_t^j(\omega) \geq 4/\varepsilon} x_t^j(\omega)\, dP(\omega|s) \geq \frac{4}{\varepsilon} P(x_t^j \geq 4/\varepsilon | s),$$

which implies that

$$P(x_t^j \geq \frac{4}{\varepsilon} | s) \leq \frac{\varepsilon}{4}.$$

Moreover, let $\bar{v}$ be the upper bound on the agents' lifetime utility coming from the boundedness of the utility function. Choose $a \in R$ such that

$$\frac{\bar{v} - U(x, \delta)}{\bar{v} - a} \leq \frac{\varepsilon}{8},$$

for all possible initial values of x and δ. Such an a exists because $U(x, \delta) > -\infty$ for all possible initial values of x and δ, as initial endowments are strictly positive by Assumption 3 and the set of initial values is finite. Then notice that $U(x, \delta) \leq E\{v_t | \mathcal{F}_0\}$ for all initial values of x and δ, because an agent always has the option to refuse any trade. Moreover

$$E\{v_t | \mathcal{F}_0\} \leq P(v_t < a | \mathcal{F}_0)\, a + (1 - P(v_t < a | \mathcal{F}_0))\, \bar{v}.$$

Combining these inequalities gives

$$P(v_t < a | \mathcal{F}_0) \leq \frac{\bar{v} - U(x, \delta)}{\bar{v} - a} \leq \frac{\varepsilon}{8}.$$

Taking unconditional expectations shows that $P(v_t < a) \leq \varepsilon/8$. Since $P(s) = 1/2$ it follows that

$P(v_t < a|s) \leq \varepsilon/4$. Next, let us use the inequality

$$\begin{aligned}
&P\Big(v_t \geq a \text{ and } |u_t - v_t| \leq \eta \text{ and } x_t^j \leq 8/\varepsilon \text{ for } j = 1, 2|s\Big) \geq \\
&1 - P(v_t < a|s) - P(|u_t - v_t| > \eta|s) - \sum_{j=1}^{2} P\left(x_t^j < \frac{4}{\varepsilon}|s\right) \geq \\
&1 - \frac{\varepsilon}{4} - \frac{\varepsilon}{4} - 2\frac{\varepsilon}{4} = 1 - \varepsilon. \qquad (12)
\end{aligned}$$

Define the set $\tilde{X} = \{x : U(x, \delta) \geq a - \eta \text{ for some } \delta \in [0, 1]\}$. We want to show that $\tilde{X}$ is closed. To see this, note that $\tilde{X} = \tilde{X}^1 \cup \tilde{X}^2$ where $\tilde{X}^j \equiv \left\{x : U(x, \delta) \geq a - \eta \text{ for some } \delta \in [0, 1], x^j \geq x^{-j}\right\}$. Consider $\tilde{X}^1$. Observe that

$$\tilde{X}^1 = \left\{x : U(x, 1) \geq a - \eta,\ x^1 \geq x^2\right\}. \qquad (13)$$

To see that this is true, suppose $x \in \tilde{X}^1$. This implies that for some δ $\pi(\delta)u(x^1)+(1-\pi(\delta))u(x^2) \geq a-\eta$. Since $x^1 \geq x^2$ and $\pi(\delta)$ is increasing, it also must be true that $\pi(1)u(x^1) + (1-\pi(1))u(x^2) \geq a-\eta$. The set on the right-hand side of (13) is closed by continuity of U, therefore, $\tilde{X}^1$ is closed. Analogously, $\tilde{X}^2$ is closed, and therefore, $\tilde{X}$ is closed. Then we can define

$$X = \tilde{X} \cap (0, 8/\varepsilon]^2.$$

Notice that $x \notin \tilde{X}$ if $x^j = 0$ for some j, and $(0, 4/\varepsilon]^2$ is bounded. Therefore, X is compact and lies in the interior of R_+^2. Moreover

$$\begin{aligned}
&\Big\{\omega : v_t(\omega) \geq a \text{ and } |u_t(\omega) - v_t(\omega)| \leq \eta \text{ and } x_t^j(\omega) \leq 4/\varepsilon \text{ for } j = 1, 2\Big\} \\
&\subset \{\omega : x_t(\omega) \in X\},
\end{aligned}$$

given that $v_t(\omega) \geq a$ and $|u_t(\omega) - v_t(\omega)| \leq \eta$ imply $u_t(\omega) \geq a - \eta$. Therefore, by (12), the set X satisfies the desired inequality $P(x_t \in X|s) \geq 1 - \varepsilon$. ■

The following is a basic probability result which will be useful throughout the appendix.

Lemma 5 *Given any $s \in \{s_1, s_2\}$, suppose there are two sets A and B, two scalars $\lambda, \eta > 0$ and a period T, such that $P((x_t, \delta_t) \in A|s) > \lambda$ and $P((x_t, \delta_t) \in B|s) > 1 - \eta$ for all $t \geq T$. Then, $P((x_t, \delta_t) \in A \cap B|s) > \lambda - \eta$ for all $t \geq T$.*

Proof. The probability $P((x_t, \delta_t) \in A|s)$ can be decomposed as

$$P((x_t, \delta_t) \in A|s) = P((x_t, \delta_t) \in A \cap B|s) + P((x_t, \delta_t) \in A \cap B^c|s),$$

where B^c is the complement of B. Moreover,

$$P((x_t, \delta_t) \in A \cap B^c|s) \leq P((x_t, \delta_t) \in B^c|s) < \eta.$$

Therefore,

$$P\left((x_t,\delta_t)\in A\cap B|s\right)=P\left((x_t,\delta_t)\in A|s\right)-P\left((x_t,\delta_t)\in A\cap B^c|s\right)>\lambda-\eta.$$

■

8.2 Proof of Proposition 1

This section is structured as follows. Lemma 6 shows that if two informed agents, with endowments in a compact set X, have different marginal rates of substitutions, they can trade and achieve a gain of at least $\Delta>0$. Then the proof of the proposition uses this result to show that in the long run the marginal rates of substitution of informed agents converge. The proof is by contradiction: if in the long run there are two groups of informed agents, of positive mass, with different marginal rates of substitution, they can improve their utility by trading with each other.

Lemma 6 *Let X be a compact set which lies in the interior of R^2_+ and ζ a positive scalar. Let the belief $\bar{\delta}\in\{0,1\}$. Suppose there are at least two endowments x and $\tilde{x}$ in X, such that the marginal rates of substitutions are sufficiently different: $\mathcal{M}(x,\bar{\delta})-\mathcal{M}(\tilde{x},\bar{\delta})\geq\zeta$. Then, there is a $\Delta>0$ with the following property: for any endowment $x\in X$ there is a trade z^* (which depends on x) such that it improves the utility of the agent with the endowment x:*

$$U(x-z^*,\bar{\delta})-U(x,\bar{\delta})\geq\Delta, \tag{14}$$

and improves utility of an agent with the endowment $\tilde{x}$:

$$U(\tilde{x}+z^*,\bar{\delta})-U(\tilde{x},\bar{\delta})\geq\Delta, \tag{15}$$

for all $\tilde{x}\in X$ such that $\mathcal{M}(\tilde{x},\bar{\delta})\leq\mathcal{M}(x,\bar{\delta})-\zeta$.

Proof. The lemma proceeds as follows. First, we fix an endowment $x\in X$ and find a trade z^* that increases the utility both of the agent $(x,\bar{\delta})$ and of all the agents $(\tilde{x},\bar{\delta})$ with marginal rate of substitution smaller than $\mathcal{M}(x,\bar{\delta})-\zeta$. At the same time, we derive a minimal utility gain $\Delta\left(x\right)$ for all the agents involved. Finally, we let x vary in X and define the lower bound on such gain, $\Delta^*=\min_{x\in X}\Delta(x)$. This gives us the desired lower bound on the utility gain for any pair of agents with $\mathcal{M}(x,\bar{\delta})-\mathcal{M}(\tilde{x},\bar{\delta})\geq\zeta$.

Step 1. Consider a given endowment x and let the marginal rate of substituion $M=\mathcal{M}(x,\bar{\delta})$. We first consider a family of trades, parametrized by ε, which make the agent with the endowment-belief combination $\left(x,\bar{\delta}\right)$ better off. In the next step, we choose ε so that the correct set of agents $(\tilde{x},\bar{\delta})$ are also better off.

Given any $\varepsilon>0$, consider the trade $z=(-\varepsilon,p\varepsilon)$ where $p=M-\zeta/2$ (if $M-\zeta<0$ we can set $p=M/2$, in this case the set of agents with $\mathcal{M}(\tilde{x},\bar{\delta})\leq M-\zeta$ is empty). Provided that $x-z>0$, the utility gains from trade $-z$, for the agent $(x, \bar{\delta})$ are, by Taylor expansion,

$$\begin{aligned}
&U(x - z, \bar{\delta}) - U(x, \bar{\delta}) = \\
&= \pi(\bar{\delta})u'(x^1)\varepsilon - (1 - \pi(\bar{\delta}))u'(x^2)p\varepsilon + \frac{1}{2}\left(\pi(\bar{\delta})u''(y^1) + (1 - \pi(\bar{\delta}))u''(y^2)p^2\right)\varepsilon^2 \\
&\geq (1 - \pi(\bar{\delta}))u'(x^2)\frac{\zeta}{2}\varepsilon + \frac{1}{2}\left(\pi(\bar{\delta})u''(y^1) + (1 - \pi(\bar{\delta}))u''(y^2)p^2\right)\varepsilon^2, \qquad (16)
\end{aligned}$$

for some $\left(y^1, y^2\right) \in \left[x^1, x^1 + \varepsilon\right] \times [x^2 - p\varepsilon, x^2]$.

For any agent with an endowment-belief pair $(\tilde{x}, \bar{\delta})$ which satisfies $\tilde{x} + z > 0$ and $\mathcal{M}(\tilde{x}, \bar{\delta}) \leq M - \zeta$, the utility gains from trade z are, by Taylor expansion,

$$\begin{aligned}
&U(\tilde{x} + z, \bar{\delta}) - U(\tilde{x}, \bar{\delta}) = \\
&= -\pi(\bar{\delta})u'(\tilde{x}^1)\varepsilon + (1 - \pi(\bar{\delta}))u'(\tilde{x}^2)p + \frac{1}{2}(\pi(\bar{\delta})u''(y^1) + (1 - \pi(\bar{\delta}))u''(y^2)p^2)\varepsilon^2 \\
&\geq (1 - \pi(\bar{\delta}))u'(\tilde{x}^2)\frac{\zeta}{2}\varepsilon + \frac{1}{2}(\pi(\bar{\delta})u''(y^1) + (1 - \pi(\bar{\delta}))u''(y^2)p^2)\varepsilon^2, \qquad (17)
\end{aligned}$$

for some $\left(y^1, y^2\right) \in \left[\tilde{x}^1 - \varepsilon, \tilde{x}^1\right] \times [\tilde{x}^2, \tilde{x}^2 + p\varepsilon]$.

Notice that since X is compact and lies in the interior of R^2_+, and given expression (16), there exists an $\bar{\varepsilon}$ such that for all $\varepsilon \in (0, \bar{\varepsilon}]$ the trade $z = (-\varepsilon, p\varepsilon)$ satisfies two properties: (i) it is feasible for agent $(x, \bar{\delta})$ and increases his utility by $U(x + z, \bar{\delta}) - U(x, \bar{\delta}) > 0$, and (ii) it is feasible for all agents with $\tilde{x} \in X$.

Step 2. Now we proceed to find an $\varepsilon^* \in (0, \bar{\varepsilon}]$ such that $z^* = (-\varepsilon^*, p\varepsilon^*)$ makes all agents $(\tilde{x}, \bar{\delta})$ better off if $\tilde{x} \in A$, where A is the compact set

$$A = \left\{\tilde{x} \in X : \mathcal{M}(\tilde{x}, \bar{\delta}) \leq M - \zeta\right\}.$$

In the process, we will find a lower bound for the utility gain of the agent $(x, \bar{\delta})$ and of all the agents with $\tilde{x} \in A$. Since A is compact the following problem is well defined and gives us a lower bound for the utility gain associated to trade $z = (-\varepsilon, p\varepsilon)$, for all the agents with endowments in A, for each $\varepsilon \in [0, \bar{\varepsilon}]$:

$$f(\varepsilon) = \min_{\substack{\tilde{x} \in A, \\ y^1 \in [\tilde{x}^1 - \varepsilon, \tilde{x}^1], \\ y^2 \in [\tilde{x}^2, \tilde{x}^2 + p\varepsilon]}} (1 - \pi(\bar{\delta}))u'(\tilde{x}^2)\frac{\zeta}{2}\varepsilon + \frac{1}{2}\left[\pi(\bar{\delta})u''(y^1) + (1 - \pi(\bar{\delta}))u''(y^2)p^2\right]\varepsilon^2.$$

Notice that at this stage $f(\varepsilon)$ may be negative for some $\varepsilon \in [0, \bar{\varepsilon}]$, as the trade may be too large and the second order terms may matter. However, there must be at least one $\varepsilon \in [0, \bar{\varepsilon}]$ such that $f(\varepsilon) > 0$. Moreover, the function $f(\varepsilon)$ is continuous by the theorem of the maximum, so we can find an $\varepsilon^* \in \arg\max_{\varepsilon \in [0, \bar{\varepsilon}]} f(\varepsilon)$ and ε^* must be strictly positive, since $f(0) = 0$. Letting $z^* = (-\varepsilon^*, p\varepsilon^*)$ it follows that

$$\Delta(x) = \min\left\{f(\varepsilon^*), U(x + z^*, \bar{\delta}) - U(x, \bar{\delta})\right\}$$

provides the desired lower bound, which satisfies both $U(x-z^*,\bar{\delta})-U(x,\bar{\delta}) \geq \Delta(x)$, and $U(\tilde{x}+z^*,\bar{\delta})-U(\tilde{x},\bar{\delta}) \geq \Delta(x)$.

Step 3. From the construction of $\Delta(x)$ in the previous steps it is possible to show that $\Delta(x)$ is continuous in both arguments. Moreover, X is compact, so $\Delta = \min_{x\in X} \Delta(x)$ is well defined, positive, and satisfies the property stated in the lemma. ■

We can now prove the main result of this section.

Proof of Proposition 1. First, we will prove that there is a sequence $\{\kappa_t(s)\}_{t=0}^{\infty}$ such that, as t increases, the marginal rates of substitution of the informed agents tend to be concentrated around $\kappa_t(s)$. That is, informed agents tend to have similar marginal rates of substitution in the long run. Next, we will complete our argument by proving that $\kappa_t(s)$ is constant over time.

Step 1. This step formally defines the contradiction.We want to prove that there is a sequence $\{\kappa_t(s)\}_{t=0}^{\infty}$, such that for all $\zeta > 0$ and $\eta > 0$ there is a T, such that

$$P\left(\left|\mathcal{M}(x_t,\delta_t)-\kappa_t(s)\right| > \zeta \mid \delta_t = \delta^I(s), s\right) < \eta \text{ for all } t \geq T.$$

Proceeding by contradiction, suppose that there are $\zeta, \eta > 0$ such that for infinitely many periods

$$P\left(\left|\mathcal{M}(x_t,\delta_t)-\kappa\right| > \zeta \mid \delta_t = \delta^I(s), s\right) > \eta \qquad \text{for all } \kappa.$$

The last expression can be rewritten as

$$P\left(\left|\mathcal{M}(x_t,\delta_t)-\kappa\right| > \zeta, \delta_t = \delta^I(s) \mid s\right) > \eta P\left(\delta_t = \delta^I(s) \mid s\right) \qquad \text{for all } \kappa. \tag{18}$$

Without loss of generality, we choose an $\eta < 2/3$.

Step 2. In this step, we want to show that when (18) holds, we can construct two sets $A, B \subset \Omega$ with the following property: in some period t there is a positive mass of informed agents in both sets, the agents in the first set have marginal rates of substitution above some constant κ_t^*, the agents in the second have marginal rates of substitution below $\kappa_t^* - \zeta$, and all these agents are sufficiently close to their long run expected utility.

We first show how to find the cutoff κ_t^*. Let X be a compact set in the interior of R_+^2, such that $P(x_t \in X \mid s) \geq 1-\alpha\eta/2$ for all $t \geq T$, for some T (such a set exists by Lemma 4). Moreover, given that there is a mass α of informed agents in period 0 and all informed agents remain informed, we have $P\left(\delta_t = \delta^I(s) \mid s\right) \geq \alpha$. Then, using Lemma 5 there must be infinitely many periods such that, for all κ, the probility that the marginal rate of substitution is away from κ is sufficiently high:

$$P\left(\left|\mathcal{M}(x_t,\delta_t)-\kappa\right| > \zeta, \delta_t = \delta^I(s), x_t \in X \mid s\right) > \eta P\left(\delta_t = \delta^I(s) \mid s\right) - \alpha\eta/2 \geq \alpha\eta/2,$$

or the probability that the marginal rate of substitution is close to κ is sufficiently low:

$$P\left(\left|\mathcal{M}(x_t,\delta_t)-\kappa\right| \leq \zeta, \delta_t = \delta^I(s), x_t \in X \mid s\right) < \alpha\eta/2 \qquad \text{for all } \kappa. \tag{19}$$

For any such period t, let $m = \inf\left\{\kappa : P\left(\mathcal{M}\left(x_t, \delta_t\right) \leq m, \delta_t = \delta^I\left(s\right), x_t \in X \mid s\right) \geq \alpha\eta/2\right\}$ and let our cutoff be $\kappa_t^* = m + \zeta/2$.

By definition, the probability that the marginal rate of subsitution is below κ_t^* is sufficiently high

$$P\left(\mathcal{M}\left(x_t, \delta_t\right) \leq \kappa_t^*, \delta_t = \delta^I\left(s\right), x_t \in X \mid s\right) \geq \alpha\eta/2 \tag{20}$$

and

$$P\left(\mathcal{M}\left(x_t, \delta_t\right) \leq \kappa_t^* - \zeta, \delta_t = \delta^I\left(s\right), x_t \in X \mid s\right) < \alpha\eta/2.$$

Moreover, (19) implies

$$P\left(\mathcal{M}\left(x_t, \delta_t\right) \in [\kappa_t^* - \zeta, \kappa_t^* + \zeta), \delta_t = \delta^I\left(s\right), x_t \in X \mid s\right) < \alpha\eta/2.$$

Combining the last two inequalities gives

$$P\left(\mathcal{M}\left(x_t, \delta_t\right) < \kappa_t^* + \zeta, \delta_t = \delta^I\left(s\right), x_t \in X \mid s\right) < \alpha\eta.$$

Then, given that $P\left(\delta_t = \delta^I\left(s\right), x_t \in X|s\right) \geq \alpha - \alpha\eta/2$, we obtain

$$\begin{aligned}
&P\left(\mathcal{M}\left(x_t, \delta_t\right) \geq \kappa_t^* + \zeta, \delta_t = \delta^I\left(s\right), x_t \in X \mid s\right) = \\
&P\left(\delta_t = \delta^I\left(s\right), x_t \in X \mid s\right) - P\left(\mathcal{M}\left(x_t, \delta_t\right) < \kappa_t^* + \zeta, \delta_t = \delta^I\left(s\right), x_t \in X \mid s\right) \geq \\
&\alpha - \alpha\eta/2 - \alpha\eta = \alpha\left(1 - \eta 3/2\right) > 0.
\end{aligned} \tag{21}$$

Given the set X and the $\zeta > 0$ defined above, let Δ be the lower bound defined in Lemma 6 and choose a $\hat{T} \geq T$ such that the utilities converged sufficiently:

$$P\left(|v_t - u_t| < \frac{1}{2}\alpha\eta\Delta, x_t \in X \mid s\right) > \frac{1}{2}\min\left\{\alpha\eta/2, \alpha\left(1 - \eta 3/2\right)\right\} \tag{22}$$

for all $t \geq \hat{T}$.

Let t be any period $t \geq \hat{T}$ such that (18) holds. We are now ready to define our two sets

$$\begin{aligned}
A &= \left\{\omega : \mathcal{M}\left(x_t, \delta_t\right) \leq \kappa_t^*, \delta_t = \delta^I\left(s\right), |v_t - u_t| < \frac{1}{2}\alpha\eta\Delta, x_t \in X\right\}, \\
B &= \left\{\omega : \mathcal{M}\left(x_t, \delta_t\right) \geq \kappa_t^* + \zeta, \delta_t = \delta^I\left(s\right), |v_t - u_t| < \frac{1}{2}\alpha\eta\Delta, x_t \in X\right\}.
\end{aligned}$$

Using (20), (21), (22) and Lemma 5 it follows that $P\left(A \mid s\right) > \alpha\eta/2$ and $P\left(B \mid s\right) > 0$. Notice that both A and B are $\mathcal{F}_t$-measurable.

Step 3. This step constructs a profitable deviation that shows that agents of sets A and B can trade with each other and increase their utility.

By Lemma 6, any agent in B can find a trade z^* improving the utility of this agent

$$U(x_t(\omega) - z^*, \delta_t(\omega)) - U(x_t(\omega), \delta_t(\omega)) \geq \Delta$$

and the utility of any agent in set A, for all $\tilde{\omega} \in A$

$$U(x_t(\tilde{\omega}) + z^*, \delta_t(\tilde{\omega})) - U(x_t(\tilde{\omega}), \delta_t(\tilde{\omega})) \geq \Delta$$

.

It remains to show that offering trade z^* is a profitable deviation for any agent who reaches set B at time t. By accepting the trade z^* the agents in A get an expected payoff

$$U(x_t(\tilde{\omega}) + z^*, \delta_t) \geq u_t(\tilde{\omega}) + \Delta > v_t(\tilde{\omega}).$$

Since $v_t(\tilde{\omega})$ is their equilibrium expected payoff, this implies that the trade is accepted by all agents in A, which implies that it is accepted with probability $\chi_t(z^*|s) > \alpha\eta/2$. Suppose an agent in B offers z^* at t and stops trading from $t+1$ on (whether or not the trade is accepted at t). The expected payoff of this strategy is

$$U(x_t(\omega), \delta_t(\omega)) + \chi_t(z^*|s)(U(x_t(\omega) - z^*, \delta_t(\omega)) - U(x_t(\omega), \delta_t(\omega))) > u_t(\omega) + \frac{1}{2}\alpha\eta\Delta > v_t(\omega).$$

Since there is a positive mass of agents in B, and $v_t(\omega)$ is their equilibrium expected payoff, this leads to a contradiction. ■

8.3 Proof of Proposition 2

Proposition 2 states that marginal rates of substitution converge for uninformed agents when $\kappa(s_1) \neq \kappa(s_2)$. The proof relies on constructing a deviation that would yield a positive utility to the uninformed agents if marginal rates of substitution failed to converge. The first part of the deviation (*learning phase*) is to offer a small trade that allows the agent to learn the signal s with sufficient precision. Once the agent's beliefs are sufficiently close to 0 or 1, we show that a positive utility gain can be achieve by trading with informed agents (*trading phase*).

Before turning to the proof, it is useful to derive three preliminary results. First, in Lemma 7 we derive a lower bound for the utility gains obtained by trading between informed agents. This lemma will be useful to construct the trading phase of our deviation. Next, Lemma 8 shows how to construct a small trade which allows the uninformed agent to learn the underlying signal s in the learning phase. Finally, Lemma 9 shows that the beliefs of uninformed agents tend to stay away from zero when the signal is s_1 and away from 1 when the signal is s_2. This lemma will be used to ensure that when the uninformed agent deviates, his learning offer allows him to get sufficiently close to the truth.

8.3.1 Gains from trade between informed agents

The following lemma shows that if, in the long run, there are a positive mass of informed agents with marginal rate of substitution below a certain threshold M' and a positive mass of informed agents with marginal rate of substitution above a threshold $M'' > M'$, the agents in the first group can make an offer at a price $p \in (M', M'')$, which is accepted with positive probability by the second group, and get a positive utility gain Δ.

This lemma is similar to Lemma 6. The difference is that in Lemma 6 we found a lower bound on the utility gain that can be achieved by any two informed agents whose marginal rate of substitution differs by at least ζ. Given any two such agents, we were allowing them to choose different trades depending on the specific values of their marginal rates of substitution. Here instead, we fix two bounds on the marginal rate of substitutions M' and M''. We then show that a minimal utility gain Δ is achieved by the *same* trade z, for any pair of agents with marginal rates of substitution one above M'' (set B), one below M' (set A). Part b) of the following lemma also shows that the trade z will be accepted by agents in the set B with a sufficiently high probability if their utilities converged.

Lemma 7 *Let X be a compact set which lies in the interior of R^2_+. Let $\bar{\delta} \in \{0,1\}$ and let M' and M'' be positive scalars, with $M' < M''$. Consider the sets of allocations*

$$\begin{aligned} A &= \left\{x : \mathcal{M}\left(x, \bar{\delta}\right) \leq M', x \in X\right\}, \\ B &= \left\{x : \mathcal{M}\left(x, \bar{\delta}\right) \geq M'', x \in X\right\}. \end{aligned}$$

For any $p \in (M', M'')$ and any $\theta > 0$, there is a $\Delta > 0$ and an $\varepsilon > 0$ such that the trade $z = \varepsilon(1, -p)$ satisfies $\|z\| < \theta$ and satisfies the following properties:

(a) $U(x - z, \delta) - U(x, \delta) \geq \Delta$ for all $(x, \delta) \in A$;

(b) if there is a $\lambda > 0$ such that $P\left(x_t \in B, \delta_t = \bar{\delta} \mid s\right) > \lambda$ for all $t \geq T$, then for any $\eta > 0$ there is a $\hat{T}$ such that $\chi_t(z|s) > \lambda - \eta$ for $t \geq \hat{T}$.

Proof. Since X is compact and lies in the interior of R^2_+, we can choose ε small enough that the trade $z = \varepsilon(1, -p)$ satisfies $\|z\| < \theta$ and the following properties hold for all $x \in A$ and all $\hat{x} \in B$:

(i) z is feasible, that is, $x - z \geq 0$ and $\hat{x} + z \geq 0$,

(ii) the expression

$$(1 - \pi(\bar{\delta}))u'(x^2)\left(p - M'\right)\varepsilon + \frac{1}{2}\left[\pi(\bar{\delta})u''(y^1) + (1 - \pi(\bar{\delta}))u''(y^2)p^2\right]\varepsilon^2$$

is positive for all $y^1 \in \left[x^1 - \varepsilon, x^1\right]$ and $y^2 \in [x^2, x^2 + p\varepsilon]$,

(iii) the expression

$$(1 - \pi(\bar{\delta}))u'(\hat{x}^2)\left(M'' - p\right)\varepsilon + \frac{1}{2}\left[\pi(\bar{\delta})u''(y^1) + (1 - \pi(\bar{\delta}))u''(y^2)p^2\right]\varepsilon^2$$

is positive for all $y^1 \in \left[\hat{x}^1, \hat{x}^1 + \varepsilon\right]$ and $y^2 \in [\hat{x}^2 - p\varepsilon, \hat{x}^2]$.

This implies that the following two problems are well defined

$$\Delta = \min_{\substack{x\in A \\ y^1\in[x^1-\varepsilon,x^1] \\ y^2\in[x^2,x^2+p\varepsilon]}} (1-\pi(\bar{\delta}))u'(x^2)\left(p-M'\right)\varepsilon + \frac{1}{2}\left[\pi(\bar{\delta})u''(y^1) + (1-\pi(\bar{\delta}))u''(y^2)p^2\right]\varepsilon^2,$$

$$\hat{\Delta} = \min_{\substack{x\in B \\ y^1\in[x^1,x^1+\varepsilon] \\ y^2\in[x^2,x^2-p\varepsilon]}} (1-\pi(\bar{\delta}))u'(x^2)\left(M''-p\right)\varepsilon + \frac{1}{2}\left[\pi(\bar{\delta})u''(y^1) + (1-\pi(\bar{\delta}))u''(y^2)p^2\right]\varepsilon^2,$$

and Δ and $\hat{\Delta}$ are positive.

Take any $x \in A$ and notice that

$$\begin{aligned} & U(x-z,\bar{\delta}) - U(x,\bar{\delta}) \\ = \;& -\pi(\bar{\delta})u'(x^1)\varepsilon + (1-\pi(\bar{\delta}))u'(x^2)p\varepsilon + \frac{1}{2}\left(\pi(\bar{\delta})u''(y^1) + (1-\pi(\bar{\delta}))u''(y^2)p^2\right)\varepsilon^2 \end{aligned}$$

for some $y^1 \in \left[x^1-\varepsilon, x^1\right]$ and $y^2 \in [x^2, x^2+p\varepsilon]$. Since $-\pi(\bar{\delta})u'(x^1)\varepsilon + (1-\pi(\bar{\delta}))u'(x^2)p\varepsilon = (1-\pi(\bar{\delta}))u'(x^2)\left(p-M'\right)\varepsilon$, this implies that

$$U(x-z,\bar{\delta}) - U(x,\bar{\delta}) \geq \Delta \text{ for all } x \in A,$$

proving part (a).

Take any $x \in B$ and notice that

$$U(x+z,\bar{\delta}) - U(x,\bar{\delta}) = \pi(\bar{\delta})u'(x^1)\varepsilon - (1-\pi(\bar{\delta}))u'(x^2)p\varepsilon + \frac{1}{2}\left(\pi(\bar{\delta})u''(y^1) + (1-\pi(\bar{\delta}))u''(y^2)p^2\right)\varepsilon^2$$

for some $y^1 \in \left[x^1, x^1+\varepsilon\right]$ and $y^2 \in [x^2-p\varepsilon, x^2]$. Since $\pi(\bar{\delta})u'(x^1)\varepsilon - (1-\pi(\bar{\delta}))u'(x^2)p\varepsilon = (1-\pi(\bar{\delta}))u'(x^2)\left(M''-p\right)\varepsilon$, this implies that

$$U(x+z,\bar{\delta}) - U(x,\bar{\delta}) \geq \hat{\Delta} \text{ for all } x \in B.$$

Now suppose that $P\left(x_t \in B, \delta_t = \bar{\delta} \mid s\right) > \lambda$ for all $t \geq T$. Using Lemmas 3 and 5, for any $\eta > 0$ we can find a $\hat{T}$ such that $P(|v_t - u_t| < \hat{\Delta}, x_t \in A, \delta_t = \bar{\delta}|s) > \lambda - \eta$ for all $t \geq \hat{T}$. Any agent with $|v_t - u_t| < \hat{\Delta}$, $x_t \in A$ and $\delta_t = \bar{\delta}$ is strictly better off accepting trade z, given that

$$U(x_t+z,\bar{\delta}) \geq U(x_t,\bar{\delta}) + \hat{\Delta} = u_t + \hat{\Delta} > v_t.$$

This shows that $\chi_t\left(z|s\right) > \lambda - \eta$, completing the proof of part (b). ■

8.3.2 Experimentation and learning when $\kappa\left(s_1\right) \neq \kappa\left(s_2\right)$

The next lemma shows that there exists a trade z that will be accepted with a sufficiently high probability in one state and rejected with a sufficiently high probability in the other state.

Lemma 8 *Suppose $\kappa(s_1) > 1 > 1/\kappa(s_2)$. For any $\theta > 0$ and any $\eta > 0$:*

(i) there is a period T and a trade z, with $\|z\| < \theta$, such that $\chi_t(z|s_1) > \alpha - \eta$, and $\chi_t(z|s_2) < \eta$ for all $t \geq T$;

(ii) there is a period T and a trade z, with $\|z\| < \theta$, such that $\chi_t(z|s_2) > \alpha - \eta$ and $\chi_t(z|s_1) < \eta$ for all $t \geq T$.

Proof. We prove part (i), the proof of part (ii) is symmetric.

Step 1. We start with the usual step — ensure that allocations of informed agents end up with sufficiently high probability in an interior compact set with given properties.

Since the marginal rate of substitution of the informed agents converge, by Lemma 1, given any $\eta > 0$, we can apply Lemmas 4 and 5, and find a compact set X in the interior of R_+^2, a positive scalar ζ, and a time $\hat{T}$ such that

$$P(\mathcal{M}(x_t, 1) \geq 1 + \zeta, x_t \in X|s_1) > \alpha - \eta/2 \text{ for all } t \geq \hat{T}. \tag{23}$$

Define the set

$$A = \{x : \mathcal{M}(x, 1) \geq 1 + \zeta, x \in X\},$$

i.e., the set of allocations in X at which the informed agents in state s_1 have marginal rate of substitution sufficiently above 1. Analogously, define the symmetric set

$$\hat{A} = \left\{x = \left(x^2, x^1\right) : \left(x^1, x^2\right) \in A\right\}.$$

In a symmetric equilibrium, the informed agents behave in a symmetric way, conditional on the signals s_1 and s_2, so that $P(x_t \in A, \delta_t = 1 \mid s_1) = P(x_t \in \hat{A}, \delta_t = 0 \mid s_2)$. Therefore, (23) implies

$$P(x_t \in \hat{A}, \delta_t = 0 \mid s_2) > \alpha - \eta/2 \text{ for all } t \geq \hat{T}, \tag{24}$$

i.e., the allocations of the informed agent if the state of the world is s_2 are in the set $\hat{A}$ with the relevant probability.

Step 2. We now construct a trade z that will be accepted by the informed agents in state s_1 with high enough probability.

Proceeding as in the proof of Lemma 7, we can find a lower bound Δ for the utility gain and choose ε small enough that the trade $z = (\varepsilon, -\varepsilon)$ satisfies $\|z\| < \theta$ and

$$x + z \geq 0 \text{ and } U(x + z, 1) - U(x, 1) \geq \Delta \text{ for all } x \in A,$$

i.e., it is interior and improves utility of the informed agents in state s_1.

The only issue we need to address in this step is to show that if the utilities of a positive mass of informed agents converges in the long run, such agents will accept the relevant trade.

Given that $P(x_t \in A, \delta_t = 1 \mid s_1) > \alpha - \eta/2$ for all $t \geq \hat{T}$, from (23), we can apply Lemmas 3 and 5 and find a T such that $P(|v_t - u_t| < \Delta, x_t \in A, \delta_t = 1|s_1) > \alpha - \eta/2 - \eta/2 = \alpha - \eta$ for all $t \geq T$. Any agent with $|v_t - u_t| < \Delta$, $x_t \in A$ and $\delta_t = 1$ is strictly better off accepting trade z, given that

$$U(x_t + z, 1) \geq U(x_t, 1) + \Delta = u_t + \Delta > v_t.$$

This shows that $\chi_t(z|s_1) > \alpha - \eta$ for all $t \geq T$.

Step 3. This step is the most difficult and important step. We need to show that we can choose T large enough that the trade will not be accepted by the informed agents in state s_2 — $\chi_t(z|s_2) < \eta$ for all $t \geq T$.

First, as in the argument in Step 2, and using symmetry we show that the opposite of trade z will be accepted by the informed agents. Take any $\hat{x} = (x^2, x^1) \in \hat{A}$, then given that $x = (x^1, x^2) \in A$ we have $x + z \geq 0$ and

$$\begin{aligned} U(\hat{x} - z, 0) - U(\hat{x}, 0) &= (1-\phi)\, u(x^2 - \varepsilon) + \phi u(x^1 + \varepsilon) - (1-\phi)\, u(x^2) - \phi u(x^1) \\ &= U(x + z, 1) - U(x, 1) \geq \Delta, \end{aligned}$$

i.e., the opposite of trade z is utility improving for the informed agents in state s_2.

Second, suppose, by contradiction, that, for some $\eta > 0$ that the probability of the offer z being accepted in state s_2 is sufficiently high: $\chi_t(z|s_2) > \eta$ for infinitely many periods. Given that $P(x_t \in \hat{A}, \delta_t = 0|s_2) > \alpha - \eta/2$ for all $t \geq \hat{T}$, we can apply Lemmas 3 and 5 and find a $\bar{T}$ such that the utility of the informed agents in the state s_2 converged sufficiently– $P(|v_t - u_t| < \eta\Delta, x_t \in \hat{A}, \delta_t = 0|s_2) > \alpha - \eta$ for all $t \geq \bar{T}$.

Pick a period $t \geq \bar{T}$ such that $\chi_t(z|s_2) > \eta$. An informed agent with $|v_t - u_t| < \eta\Delta$, $x_t \in \hat{A}$ and $\delta_t = 0$ is strictly better off making the offer z, consuming x_t if the offer is rejected, and consuming $x_t - z$ if the offer is accepted, given that

$$\begin{aligned} (1 - \chi_t(z|s_2))\, U(x_t, 0) + \chi_t(z|s_2)\, U(x_t - z, 0) &\geq \\ U(x_t, 0) + \eta\,(U(x_t - z, 0) - U(x_t, 0))\,\Delta &\geq u_t + \eta\Delta > v_t. \end{aligned}$$

Since this behavior dominates the equilibrium strategy and there is a positive mass of informed agents with $|v_t - u_t| < \Delta$, $x_t \in \hat{A}$ and $\delta_t = 0$, we have a contradiction ■

8.3.3 A bound on incorrect beliefs

The following lemma shows that conditional on the signal s_1, there is a small probability that the belief of the uninformed agent gets too close to 0. That is, the uninformed agents can only be very wrong with a small probability. This lower bound will be useful when we construct our profitable deviation in the proof of Proposition **2**.

Lemma 9 *For each $\varepsilon \in (0,1)$ the bound on the incorrect beliefs for all t is given by:*

$$P(\delta_t < \varepsilon \mid s_1) < \frac{\varepsilon}{1-\varepsilon}.$$

Proof. Since $\delta_t(\omega)$ are equilibrium beliefs, they must be consistent with Bayesian updating and must satisfy, by definition, $\delta = P(s = s_1 \mid \delta_t(\omega) = \delta)$. This implies $P(s_1 \mid \delta_t < \varepsilon) < \varepsilon$, which implies $P(s_2 \mid \delta_t < \varepsilon) > 1 - \varepsilon$ and

$$\frac{P(s_1 \mid \delta_t < \varepsilon)}{P(s_2 \mid \delta_t < \varepsilon)} < \frac{\varepsilon}{1-\varepsilon}.$$

Applying Bayes' rule

$$\frac{P(s_1 \mid \delta_t < \varepsilon)}{P(s_2 \mid \delta_t < \varepsilon)} = \frac{P(\delta_t < \varepsilon \mid s_1) P(s_1)}{P(\delta_t < \varepsilon \mid s_2) P(s_2)}.$$

Combining the last two equations, and using $P(s_1) = P(s_2) = 1/2$, gives

$$P(\delta_t < \varepsilon \mid s_1) < \frac{\varepsilon}{1-\varepsilon} P(\delta_t < \varepsilon \mid s_2) \leq \frac{\varepsilon}{1-\varepsilon},$$

where the last inequality follows from $P(\delta_t < \varepsilon \mid s_2) \leq 1$. ■

We can now turn to the main result of this section.

Proof of Proposition 2.

Suppose, by contradiction, that there exist an $s \in \{s_1, s_2\}$, a $\zeta > 0$ and a $\xi > 0$ such that

$$P\left(\left|\mathcal{M}\left(x_t, \delta^I(s)\right) - \kappa(s)\right| > \zeta \mid s\right) > \xi \tag{25}$$

for infinitely many periods. Without loss of generality, suppose this holds for $s = s_1$. The case $s = s_2$ is treated in a symmetric way. We will show that (25) is incompatible with optimality for uninformed agents. An uninformed agent can construct a strategy which leads to a strictly higher payoff than the equilibrium strategy. The profitable deviation consists of two phases.

(i) *Learning phase.* The uninformed agent makes a small offer $\hat{z}$ to learn the value of the signal s . One can think of this phase as strategic experimentation by the uninformed agents to learn the state. We use Lemma 8 to construct this part of the proof.

(ii) *Trading phase.* The uninformed agent makes an offer z^*, which may take one of the two values z^- or z^+, depending on the agent's marginal rate of substitution. This offer generates a potential gain of Δ. We use Lemma 7 to construct this part of the proof.

In the following steps, we construct this deviation in detail and then we show that such strategy indeed makes an uninformed agent strictly better off, i.e., that the costs of the learning phase overweigh the gains of the trading phase.

Step 1. This is the usual first step of the proofs in the paper, in which we ensure that endowments and beliefs converge to a set where certain conditions are satisfied. First, we define a compact set X and find a lower bound for the belief δ_t. Specifically, we proceed as follows. Choose two positive scalars ν and ε such that $\nu < \min\{\xi/2, \alpha/2\}$ and $\varepsilon/(1-\varepsilon) < \xi/4$. By Lemma 4, there is a compact set such that the allocations converged to this set – $P(x_t \in X \mid s_1) \geq 1 - \nu$ for infinitely many periods and by

Lemma 9 $P(\delta_t \geq \varepsilon \mid s_1) > 1 - \varepsilon/(1-\varepsilon)$ for all t . Then, applying Lemma 5, we can show that (25) implies that

$$P(|\mathcal{M}(x_t, 1) - \kappa(s_1)| > \zeta, \delta_t \geq \varepsilon, x_t \in X \mid s_1) > \xi - \nu - \varepsilon/(1-\varepsilon) > \xi/4 \tag{26}$$

holds for infinitely many periods, i.e., there is a sufficiently large mass of uninformed agents with marginal rate of substitution far enough from $\kappa(s_1)$ and who are sufficiently optimistic about s_1 . Proposition 1 implies that there is a T_1 such that

$$P(|\mathcal{M}(x_t, \delta_t) - \kappa(s_1)| < \zeta/2, \delta_t = 1, x_t \in X \mid s_1) > \alpha/2 \tag{27}$$

for all $t \geq T_1$, i.e., when the marginal rate of substitutions has converged to κ_1 for a sufficient mass of informed agents.

Step 2. Now we begin constructing the deviating strategy. Going backward, we first find the trade z^* offered in the trading phase, after the agent has learnt the signal s with a sufficient precision. We also define the minimum utility gain Δ achieved in the trading phase.

Consider the sets

$$\begin{aligned} A^- &= \{x : \mathcal{M}(x, 1) \leq \kappa(s_1) - \zeta, x \in X\}, \\ B^- &= \{x : \mathcal{M}(x, 1) \geq \kappa(s_1) - \zeta/2, x \in X\}. \end{aligned}$$

Applying Lemma 7, we can find a minimum utility gain $\Delta^- > 0$ and a trade z^- sufficiently small that $\bar{X} = \{\tilde{x} = x + z^- : x \in X\}$ is a non-empty compact set in the interior of R^2. By part (a) of Lemma 7, the trade z^- has the property that $U(x - z^-, 1) - U(x, 1) \geq \Delta^-$ for all $x \in A^-$. Moreover, from (27) it follows that $P(x_t \in B, \delta_t = 1 \mid s) > \alpha/2$ for all $t \geq T_1$, and from part (b) of Lemma 7 we can find a T_2^- such that $\chi_t(z^-|s_1) > \alpha/4$ for all $t \geq T_2^-$. We can do a similar construction considering the sets

$$\begin{aligned} A^+ &= \{x : \mathcal{M}(x, 1) \geq \kappa(s_1) + \zeta, x \in X\}, \\ B^+ &= \{x : \mathcal{M}(x, 1) \leq \kappa(s_1) + \zeta/2, x \in X\}, \end{aligned}$$

and find a minimum utility gain $\Delta^+ > 0$ and a trade z^+ with $\chi_t(z^+|s_1) > \alpha/4$ for all $t \geq T_2^+$. Let $\Delta = \min\{\Delta^-, \Delta^+\}$.

Step 3. We now define the trade $\hat{z}$ that is offered in the learning phase. We will also define some bounds that will be used to check that our deviation is profitable ex ante.

First, we need to choose the value of the positive scalar η, which represents the maximum probability that trade $\hat{z}$ is accepted in the wrong state s_2. That is, η captures the accuracy of the signal obtained by offering $\hat{z}$. Let us define an upper bound for the utility loss of an agent making offer z^+ or z^- in the event that the signal turns out to be s_2:

$$Q \equiv \max_{\substack{x \in X \\ z \in \{z^+, z^-\}}} U(x, 0) - U(x + z, 0).$$

If $Q > 0$ choose a positive scalar η such that

$$\frac{\eta}{\alpha - \eta} < \frac{\gamma\alpha}{24} \frac{\varepsilon}{1 - \varepsilon} \frac{\Delta}{Q}.$$

If $Q \leq 0$ then choose any arbitrary $\eta > 0$.

Second, we need to choose the positive scalar $\theta > 0$. Notice that we want to ensure that the trade $\hat{z}$ is small enough that, by itself , it only causes a negligible change in utility. Therefore, we choose $\theta > 0$ such that that if $x' \in \bar{X}$ and $\|x'' - x'\| < \theta$ then $x'' > 0$ and

$$\left| U\left(x'', \delta\right) - U\left(x', \delta\right) \right| < \lambda, \tag{28}$$

for $\delta \in \{0, 1\}$, where

$$\lambda = \varepsilon \frac{\gamma\alpha}{24} (\alpha - \eta) \Delta.$$

We are now ready for the key observation of this step. Given η and θ, apply Lemma 8 and find a period T_3 and a trade $\hat{z}$, with $\|\hat{z}\| < \theta$, such that

$$\chi_t(\hat{z}|s_1) > \alpha - \eta, \quad \chi_t(\hat{z}|s_2) < \eta, \text{ for all } t \geq T_3. \tag{29}$$

Finally, we want to ensure that the expected utility of the uninformed agent has converged at the time of the profitable deviation. To do so we define

$$\hat{\Delta} = \varepsilon \frac{\gamma\alpha}{24} (\alpha - \eta) \Delta.$$

Then, applying Lemmas 3 and 5, from (26) it follows that

$$P\left(|\mathcal{M}(x_t, 1) - \kappa(s_1)| > \zeta, \delta_t \geq \varepsilon, |v_t - u_t| < \hat{\Delta}, x_t \in X \mid s_1 \right) > \xi/8 \tag{30}$$

holds for infinitely many periods.

Step 4. In this step we construct a profitable deviation for uninformed agents.

Let T a period larger than T_1, T_2^-, T_2^+ and T_3, such that (30) holds. The agent follows the equilibrium strategy σ before T. In period T, provided the game does not end, he deviates if all the following conditions hold: he is selected to make an offer in T; the difference in the MRS if he knew the state is significant ,$|\mathcal{M}(x_T, 1) - \kappa(s_1)| > \zeta$; the belief is interior $\delta_T \geq \varepsilon$; the utilities have converged $|v_T - u_T| < \hat{\Delta}$, and $x_T \in X$. If any of these conditions fails, he keeps on playing σ. By construction, we know that a deviation happens with a positive probability. The deviation is as follows: he offers $\hat{z}$ in period T; if his offer is accepted and he is selected propose the offer in period $T + 1$, he makes the offer z^*, where

$z^* = z^-$ if $\mathcal{M}(x_T, 1) < \kappa(s_1) - \zeta$ and $z^* = z^+$ if $\mathcal{M}(x_T, 1) > \kappa(s_1) + \zeta$. In all the other cases he stops trading and waits for the game to end to consume his endowment.

Under this strategy, three possible outcomes arise:

1. The agent makes the offers $\hat{z}$ and z^* and both are accepted. In this case, the final endowment is $x_T + \hat{z} + z^*$. The probability of this outcome, conditional on the signal s, is

$$\Pi_1(s) = \frac{1}{2}\gamma\chi_T(\hat{z}|s)\,\chi_{T+1}(z^*|s),$$

because in period $T+1$ the game must not end and the agent must be selected as the proposer.

2. Offer $\hat{z}$ is accepted but either the agent is not selected to propose an offer in $T+1$, or his offer z^* is rejected. In this case, the final endowment is $x_T + \hat{z}$. The probability of this outcome, conditional on the signal s, is

$$\Pi_2(s) = \frac{1}{2}\gamma\chi_T(\hat{z}|s)\left(1 - \chi_{T+1}(z^*|s)\right) + \left(1 - \frac{1}{2}\gamma\right)\chi_T(\hat{z}|s).$$

3. The offer $\hat{z}$ is rejected and the agent consumes x_T, $\Pi_3(s) = 1 - \Pi_1(s_1) - \Pi_2(s_1)$.

Since ex ante the agent assigns probabilities δ_T and $1 - \delta_T$ to the signals s_1 and s_2, the expected payoff of an agent at the moment of deviating is

$$\begin{aligned}\hat{U} &= \delta_T\left[\Pi_1(s_1)\,U(x_T + \hat{z} + z^*, 1) + \Pi_2(s_1)\,U(x_T + \hat{z}, 1) + \Pi_3(s_1)\,U(x_T, 1)\right] + \\ &+ (1-\delta_T)\left[\Pi_1(s_2)\,U(x_T + \hat{z} + z^*, 1) + \Pi_2(s_2)\,U(x_T + \hat{z}, 1) + \Pi_3(s_2)\,U(x_T, 1)\right].\end{aligned}$$

Rearranging, using the linearity of $U(x, \delta)$ in δ, and using $u_T = U(x_T, \delta_T)$, we obtain

$$\begin{aligned}\hat{U} &= u_T + \delta_T\left[\Pi_1(s_1)\left(U(x_T + \hat{z} + z^*, 1) - U(x_T, 1)\right) + \Pi_2(s_1)\left(U(x_T + \hat{z}, 1) - U(x_T, 1)\right)\right] + \\ &+ (1-\delta_T)\left[\Pi_1(s_2)\left(U(x_T + \hat{z} + z^*, 0) - U(x_T, 0)\right) + \Pi_2(s_2)\left(U(x_T + \hat{z}, 0) - U(x_T, 0)\right)\right].\end{aligned} \tag{31}$$

By construction, the probabilities satisfy a number of inequalities. First,

$$\Pi_1(s_1) > \frac{1}{8}\gamma\alpha(\alpha - \eta)$$

follows because $\chi_T(\hat{z}|s_2) > \alpha - \eta$, from (29), and $\chi_T(z^*|s_1) > \alpha/4$, from Step 2. Second,

$$\Pi_1(s_2) + \Pi_2(s_2) = \frac{1}{2}\gamma\chi_T(\hat{z}|s_2) + \left(1 - \frac{1}{2}\gamma\right)\chi_T(\hat{z}|s_2) < \eta,$$

follows because $\chi_T(\hat{z}|s_2) < \eta$, from (29). Third,

$$\delta_T\left(\Pi_1(s_1) + \Pi_2(s_1)\right) + (1-\delta_T)\left(\Pi_1(s_2) + \Pi_2(s_2)\right) \le 1,$$

follows immediately. Finally $\delta_T \geq \varepsilon$ and $1-\delta_T < 1-\varepsilon$, by construction. Moreover, the utility changes are bounded as follows

$$\begin{aligned}
U\left(x_T+\hat{z}+z^*,1\right)-U\left(x_T,1\right) &\geq \Delta-\lambda, \\
U\left(x_T+\hat{z},1\right)-U\left(x_T,1\right) &\geq -\lambda, \\
U\left(x_T+\hat{z}+z^*,0\right)-U\left(x_T,0\right) &\geq -Q-\lambda, \\
U\left(x_T+\hat{z},0\right)-U\left(x_T,0\right) &\geq -\lambda,
\end{aligned}$$

which exploit the definition of Q, and property (28) —both in Step 3—, and the triangle inequality. Applying these bounds for probabilities and utility changes to the right-hand side of (31), we obtain

$$\begin{aligned}
\hat{U} &\geq u_T+\varepsilon\frac{1}{8}\gamma\alpha\left(\alpha-\eta\right)\left(\Delta-\lambda\right)-\varepsilon\left(1-\frac{1}{8}\gamma\alpha\left(\alpha-\eta\right)\right)\lambda-\left(1-\varepsilon\right)\eta\left(Q+\lambda\right)= \\
&= u_T+\varepsilon\frac{1}{8}\gamma\alpha\left(\alpha-\eta\right)\Delta-\left(1-\varepsilon\right)\eta Q-\left(\varepsilon+\left(1-\varepsilon\right)\eta\right)\lambda> \\
&> u_T+\varepsilon\frac{1}{24}\gamma\alpha\left(\alpha-\eta\right)\Delta>v_T.
\end{aligned}$$

The definitions of η and λ, in Step 3, ensure that $\left(1-\varepsilon\right)\eta Q$ and $\left(\varepsilon+\left(1-\varepsilon\right)\eta\right)\lambda$ are each smaller than one third of the utility gain $\varepsilon\gamma\alpha\left(\alpha-\eta\right)\Delta/8$, which ensures that the second inequality in this expression holds. The definition of $\hat{\Delta}$, also in Step 3, and the fact that $|v_T-u_T|<\hat{\Delta}$ ensure that the last inequality holds. This shows that the deviation is profitable for the uninformed agent. Since this happens with positive probability, the strategy σ is suboptimal and we have a contradiction. ■

8.3.4 Proof of Proposition 3

As in the proof of Proposition 2, we need to find trades that allow the agent to learn the true state. The next lemma shows that the uninformed agent can always find a trade that reveals *some* information about the state s, that is, a trade z such that the probabilities of acceptance are sufficiently different in the two states:

$$\left|\chi_t\left(z|s_2\right)-\chi_t\left(z|s_1\right)\right|>\eta.$$

The difference with Lemma 8, is that there the trade could be chosen to reveal almost perfect information about s_1, as we could make the probability of acceptance arbitrarily close to zero in state s_2. So in that case, one round of experimentation was enough to achieve information on s_1 with any degree of precision. Here instead, the agent may need to experiment for several rounds before being sufficiently well informed. This will affect the proof of convergence below.

Lemma 10 *Suppose $\kappa(s_1)=\kappa(s_2)=1$. Suppose there is a non-vanishing mass of agents with marginal rate of substitution different from 1, that is, for some $s\in\{s_1,s_2\}$ there is a $\zeta>0$ and a $\xi>0$ such that*

$$P\left(\left|\mathcal{M}\left(x_t,\delta_t\right)-1\right|>\zeta|s\right)>\xi$$

for infinitely many periods. Then, for any $\theta > 0$ there is an $\eta > 0$ such that, for infinitely many periods there is a trade z with $\|z\| < \theta$ and

$$|\chi_t(z|s_2) - \chi_t(z|s_1)| > \eta.$$

Proof. We proceed by contradiction, supposing that the condition on $|\chi_t(z|s_1) - \chi_t(z|s_2)|$ stated in the lemma is not satisfied and showing that then there is a profitable deviation for a positive mass of uninformed agents.

Namely, suppose, that there is a $\theta^* > 0$ such that for all $\eta > 0$ there is a T such that $|\chi_t(z|s_2) - \chi_t(z|s_1)| < \eta$ for all $t \geq T$ and all z such that $\|z\| < \theta^*$. This means that, after some period T, all the trades in a ball of radius θ^* are accepted with similar probabilities.

The intuition for the proof is as follows. In the long run, an uninformed agent can make an offer at a price near 1 and expect the offer to be accepted with similar probabilities in the two states. The fact that the offer is accepted reveals little information to the uninformed agent. Then the uninformed agent's gains from trade can be approximately evaluated using his ex ante marginal rate of substitution $\mathcal{M}(x_t, \delta_t)$. Therefore, for agents with $\mathcal{M}(x_t, \delta_t)$ different from 1, a profitable deviation is possible. The rest of the proof makes this argument formal and is divided in two steps.

Step 1. This step deal with some preliminary issues. The hypothesis of the lemma implies that, for some s, either there is a $\zeta > 0$ and a $\xi > 0$ such that the marginal rate of substitution for the uninformed agent is sufficiently different from 1 for infinitely many periods:

$$P(\mathcal{M}(x_t, \delta_t) < 1 - \zeta | s) > \xi,$$

or

$$P(\mathcal{M}(x_t, \delta_t) > 1 + \zeta | s) > \xi.$$

Let us focus on the first case and suppose that it holds for $s = s_1$. All the other possible cases are proved in a similar way.

Now, we have to do our usual step to show that allocations end up in the interior compact set with relevant properties. By Proposition 1, it is possible to find a compact set X such that for any $\eta > 0$ there is a time T such that the marginal rates of substitution for informed agents converge to 1:

$$P(\mathcal{M}(x_t, \delta_t) \geq 1 - \zeta/2, \delta_t = 1, x \in X | s_1) > \alpha - \eta \text{ for all } t \geq T.$$

Choose the compact set $\hat{X}$ in the interior of R_+^2 such that the allocations of uninformed agents with the marginal rate of substitution different from 1 converge there with a sufficiently high probability, i.e., such that $P(\mathcal{M}(x_t, \delta_t) < 1 - \zeta, x_t \in \hat{X} | s_1) > \xi/2$. By Lemma 3, we can find a $\hat{T} \geq T$ such that additionally, the utilities of these agents also converge

$$P(|u_t - v_t| < \lambda\Delta/2, \mathcal{M}(x_t, \delta_t) < 1 - \zeta, x_t \in \hat{X} | s_1) > \xi/4$$

for all $t \geq \hat{T}$. We are now done with the preliminaries about convergence.

Step 2. We construct the trade z^* for the uninformed agent's profitable deviation, which is described in the next step. By Lemma 7, we can find a trade z^* that will be accepted by informed agents with sufficiently high probability if the state is s_1. Formally, there is a $z^* = \varepsilon \left(1, -\left(1 - \zeta/3\right)\right)$, with $\|z^*\| < \theta^*$, such that $\chi_t \left(z^* | s_1\right) \geq \chi_t^I \left(z^* | s_1\right) > \lambda$ for infinitely many periods for some $\lambda > 0$. Then find a $\Delta > 0$ such that if this trade indeed improves utility of the uninformed agent

$$U \left(x + z^*, \delta\right) - U \left(x, \delta\right) \geq \Delta$$

whenever $\mathcal{M} \left(x, \delta\right) < 1 - \zeta$ and $x \in \hat{X}$ (such a Δ exists by arguments similar to those in Lemma 7 and the compactness of $\hat{X}$).

Step 3. This step shows that the offer z^* constructed by the uninformed agent indeed leads to an increase in his expected utility.

Let $Q = \max_{x \in \hat{X}} \left\{U \left(x + z^*, 0\right) - U \left(x, 0\right)\right\}$, i.e. it is the highest gain of such a trade z^* that an informed agent can get if the state is s_2. Choose $\eta < \lambda\Delta/(2Q)$ and find a $t \geq \hat{T}$ such that $\left|\chi_t \left(z | s_2\right) - \chi_t \left(z | s_1\right)\right| < \eta$ for all z such that $\|z\| < \theta^*$, i.e. the trade is accepted with approximately the same probability in both states.

Now take any uninformed agent whose utility converged, $|u_t - v_t| < \lambda\Delta/2$, whose marginal rate of substitution is lower than 1, $\mathcal{M} \left(x_t, \delta_t\right) < 1 - \zeta$, and allocations are in the relevant set $x_t \in \hat{X}$ who is selected to be the proposer at time t.

Consider the following deviation. Suppose he makes the offer z^* and, after the current round of trade and consumes his endowment. His expected utility is:

$$\begin{aligned}
\hat{U} &= \delta_t \left[\left(1 - \chi_t \left(z^* | s_1\right)\right) U \left(x_t, 1\right) + \chi_t \left(z^* | s_1\right) U \left(x_t + z^*, 1\right)\right] + \\
&\quad \left(1 - \delta_t\right) \left[\left(1 - \chi_t \left(z^* | s_2\right)\right) U \left(x_t, 0\right) + \chi_t \left(z^* | s_2\right) U \left(x_t + z^*, 0\right)\right],
\end{aligned}$$

which can be rewritten, exploiting the linearity of $U \left(x, \delta\right)$ in δ, as

$$\begin{aligned}
\hat{U} &= U \left(x_t, \delta_t\right) + \chi_t \left(z^* | s_1\right) \left(U \left(x_t + z^*, \delta_t\right) - U \left(x_t, \delta_t\right)\right) + \\
&\quad \left(1 - \delta_t\right) \left(\chi_t \left(z^* | s_2\right) - \chi_t \left(z^* | s_1\right)\right) \left(U \left(x_t + z^*, 0\right) - U \left(x_t, 0\right)\right).
\end{aligned}$$

Since $\mathcal{M} \left(x_t, \delta_t\right) < 1 - \zeta$ and $x_t \in \hat{X}$ by construction $U \left(x_t + z^*, \delta_t\right) - U \left(x_t, \delta_t\right) \geq \Delta$. Moreover, the offer is accepted with high enough probability in the state 1, $\chi_t \left(z | s_1\right) > \lambda$, and the probability of having the offer accepted in the other state is approximately the same $\left|\chi_t \left(z | s_2\right) - \chi_t \left(z | s_1\right)\right| < \eta$, and the utility gain (or loss) in the state s_2 is bounded by $U \left(x_t + z^*, 0\right) - U \left(x, 0\right) \geq Q$. Then the utility from deviating at time t satisfies

$$\hat{U} \geq u_t + \lambda\Delta - \eta Q > u_t + \lambda\Delta/2,$$

and since $v_t < u_t + \lambda\Delta/2$ this is a profitable deviation for this agent. Therefore, since, by construction, there is a positive mass of agents with $|u_t - v_t| < \lambda\Delta/2$, $\mathcal{M} \left(x_t, \delta_t\right) < 1 - \zeta$ and $x_t \in \hat{X}$, we have a contradiction. ■

Proof of Proposition 3.

To be written... ■

8.4 Proof of Proposition 4

The following lemma shows that at any time t we can start from the equilibrium joint distribution of endowment and beliefs, $\Gamma_t(.|s_1)$, and eliminate symmetric masses of agents with $\delta < 1/2$ and $\delta \geq 1/2$. By this process, we end up with a distribution of endowment and beliefs $\tilde{\Gamma}_t(.|s_1)$, where every agent has $\delta \geq 1/2$, and the average endowments of goods 1 and 2 are equal.

Lemma 11 *Let $W = R_+^2 \times [1/2, 1]$ be the set of endowment-belief pairs such that $\delta \geq 1/2$. For any Borel set $F \subseteq W$ let F^C denote the set $\left\{\left(x^1, x^2, \delta\right) : \left(x^2, x^1, 1-\delta\right) \in F\right\}$ and let*

$$\tilde{\Gamma}_t(F|s_1) \equiv P((x_t, \delta_t) \in F|s_1) - P\left((x_t, \delta_t) \in F^C|s_1\right).$$

$\tilde{\Gamma}_t$ is a measure on W with the following property

$$\int_W \left(x^2 - x^1\right) d\tilde{\Gamma}_t = 0 \text{ for } j = 1, 2.$$

Proof. Notice that by construction $\delta_t \geq 1/2$ if $(x_t, \delta_t) \in F$. Given the definition of δ_t and given that x_t is $\mathcal{F}_t$-measurable, we then have $P(s_1|F) \geq 1/2$. Moreover, Bayes' rule implies that

$$\frac{P(s_1|F)}{P(s_2|F)} = \frac{P(F|s_1)P(s_1)}{P(F|s_2)P(s_2)},$$

which, together with $P(s_1|F) \geq 1/2$ and $P(s_1) = P(s_2)$ implies $P(F|s_1) \geq P(F|s_2)$ (this can hold with equality only if all the points in F, except a set of zero measure under both $P(.|s_1)$ and $P(.|s_2)$, have $\delta = 1/2$). By symmetry, we have $P(F|s_2) = P\left(F^C|s_1\right)$, and thus $P(F|s_1) \geq P\left(F^C|s_1\right)$. This shows that $\tilde{\Gamma}_t(F|s_1) \geq 0$ for all $F \subseteq W$.

Recall that $\Gamma_t(F|s_1) = P((x_t, \delta_t) \in F|s_1)$ and let $\Gamma_t^C(.|s_1)$ be defined as

$$\Gamma_t^C(F|s_1) = \Gamma_t\left(F^C|s_1\right).$$

We will use Γ_t, Γ_t^C and $\tilde{\Gamma}_t$ as shorthand for $\Gamma_t(.|s_1), \Gamma_t^C(.|s_1)$ and $\tilde{\Gamma}_t(.|s_1)$. Using market clearing we have

$$\int_{x^2>x^1} \left(x^2 - x^1\right) d\Gamma_t + \int_{x^2=x^1} \left(x^2 - x^1\right) d\Gamma_t - \int_{x^2<x^1} \left(x^1 - x^2\right) d\Gamma_t = \int \left(x^2 - x^1\right) d\Gamma_t = 0,$$

and the middle term in the first expression is zero. By construction $\int_{x^2<x^1}\left(x^1-x^2\right)d\Gamma_t=\int_{x^1<x^2}\left(x^2-x^1\right)d\Gamma_t^C$, so, substituting, we have

$$\int_{x^2>x^1}\left(x^2-x^1\right)\left(d\Gamma_t-d\Gamma_t^C\right)=0.$$

Decomposing the term on the right-hand side gives

$$\int_{\substack{x^2>x^1\\ \delta\geq 1/2}}\left(x^2-x^1\right)\left(d\Gamma_t-d\Gamma_t^C\right)+\int_{\substack{x^2>x^1\\ \delta<1/2}}\left(x^2-x^1\right)\left(d\Gamma_t-d\Gamma_t^C\right)=0.$$

Again, by construction, $\int_{\substack{x^2>x^1\\ \delta<1/2}}\left(x^2-x^1\right)\left(d\Gamma_t-d\Gamma_t^C\right)=\int_{\substack{x^1\geq x^2\\ \delta\geq 1/2}}\left(x^1-x^2\right)\left(d\Gamma_t^C-d\Gamma_t\right)$ (where the cases $x^1=x^2$ and $\delta=1/2$ are allowed because in the first case $x^1-x^2=0$, in the second case $d\Gamma_t^C-d\Gamma_t=0$). We conclude that

$$\int_W\left(x^2-x^1\right)\left(d\Gamma_t-d\Gamma_t^C\right)=\int_W\left(x^2-x^1\right)d\tilde{\Gamma}_t=0.$$

■

Proof of Proposition 4. We sketch the proof here. Write

$$\frac{u'\left(x_t^1\right)}{u'\left(x_t^2\right)}=\frac{1-\pi\left(\delta_t\right)}{\pi\left(\delta_t\right)}\mathcal{M}(x_t,\delta_t).$$

If $\delta_t\geq 1/2$ this implies

$$\frac{u'\left(x_t^1\right)}{u'\left(x_t^2\right)}\leq\mathcal{M}(x_t,\delta_t).$$

Using the fact that $\mathcal{M}(x_t,\delta_t)$ is approximately 1 for most agents, we prove that if $\delta_t\geq 1/2$ either $x_t^1>x_t^2$ or $x_t^1\approx x_t^2$. Moreover, there is a positive mass of informed with $x_t^1>x_t^2$. By the previous lemma we can integrate only over agents with $\delta_t\geq 1/2$ (using the measure $\tilde{\Gamma}_t$) and since we are integrating over agents who either have $x_t^1>x_t^2$ or $x_t^1\approx x_t^2$ we get a contradiction because market clearing requires $\int_W\left(x^1-x^2\right)d\tilde{\Gamma}_t=0$. ■

9 Computational Appendix

This appendix describes computational algorithms we used to compute numerical examples in Section ??.

In a Markov equilibrium considered throughout the paper, agent's strategy in period t depends on his allocation of assets inherited from the previous period, x_{t-1}, his beliefs about the probability of signal s_1, δ_{t-1}, and the distribution of beliefs and endowments of other agents $\Gamma_t(\cdot|s)$ for $s=\{s_1,s_2\}$. Notice that an individual agent cannot affect the distribution $\{\Gamma_t(\cdot|s)\}_{t=0}^{\infty}$ since agent's actions are observable only to a measure zero of agents. Therefore, each agent treats the sequence $\{\Gamma_t(\cdot|s)\}_{t=0}^{\infty}$ as given, and the dependence on that sequence can be summarize by the calendar time t, so that the state of each agent is (x,δ,t).

At the beginning of period t, an agent has assets x_{t-1} and beliefs δ_{t-1} and chooses his optimal strategy σ_t to maximize the payoff $W(x_{t-1}, \delta_{t-1}, t)$:

$$\begin{aligned} W(x_{t-1}, \delta_{t-1}, t) &= \max_{\sigma_t} (1-\gamma) E\left\{ U(x_t(\sigma_t), \delta_t(\sigma_t)) \middle| \Pr(s = s_1) = \delta_{t-1} \right\} \\ &\quad + \gamma E\left\{ W(x_t(\sigma_t), \delta_t(\sigma_t), t+1) \middle| \Pr(s = s_1) = \delta_{t-1} \right\}. \end{aligned}$$

An implication of the expression above is that agent's best response strategy $\sigma^* = \{\sigma_t^*\}_{t=1}^{\infty}$ in the infinitely repeated game consists of a sequence of the best responses σ_t^* in a static game where agent's payoff is given by $(1-\gamma)\, U(\cdot,\cdot) + \gamma W(\cdot,\cdot,t+1)$. Therefore, if a sequence of payoffs $\{W(\cdot,\cdot,t)\}_{t=1}^{\infty}$, one can find equilibrium strategies of agents by the following recursive procedure:

1. Start with the initial distribution $\Gamma_0(\cdot|s)$ and compute a static Bayesian Nash equilibrium of this game with payoffs $(1-\gamma)\, U(\cdot,\cdot) + \gamma W(\cdot,\cdot,1)$;

2. Use equilibrium strategies to compute the distribution in the next period, $\Gamma_1(\cdot|s)$; compute static Bayesian Nash equilibrium for the period $t = 1$;

3. Repeat the above procedure for periods $t = 2, 3, ...$

The two crucial ingredients of this procedure are: (1) finding the sequence of payoffs $\{W(\cdot,\cdot,t)\}_{t=1}^{\infty}$; and (2) finding an equilibrium in a static game with an arbitrary distribution of beliefs and endowments $\Gamma(x, \delta|s)$ and payoffs $((1-\gamma)\, U + \gamma W)$.

Now we describe a general procedure to compute an equilibrium in our games. Then, we discuss some further simplifications we used for computations in Section **??**.

For computational purposes, we discretize the state space and the set of offers that agents can make as follows. We fix a grid size (the step of the grid) for the offers to be h_z and for the beliefs to be h_δ. We set the bound for the size of the maximal allowed offer as $\bar{z}$, and the set of allowable offers consists is given by $\mathbf{Z} = Z \times Z$, with $Z \equiv \{\pm n h_z : |n h_z| \leq \bar{z},\ n \in \mathbb{N}\}$, where $\mathbb{N}$ is a set of natural numbers. Similarly, allocations of agents take values on a set $\mathbf{X} = X \times X$, with $X \equiv \{\pm n h_z : |n h_z| \leq \bar{x},\ n \in \mathbb{N}\}$ where $\bar{x}$ is a bound on agent's allocations. Agent's beliefs take values on a set $\Delta \equiv \{0, h_\delta, 2h_\delta, ..., 1\}$.

9.1 Finding an equilibrium in a static game

The first step is to compute an equilibrium in a static, one shot game for some distribution $\Gamma : \mathbf{X} \times \Delta \to [0,1]$ and payoffs $W : \mathbf{X} \times \Delta \to \mathbb{R}$. For this purpose we adopt the algorithm of Fudenberg and Levine (1995) to our Bayesian game. This algorithm computes an approximate equilibrium for a static game, where a degree of approximation depends on a parameter κ. The algorithm has a property as $\kappa \to \infty$ the equilibrium strategies in the approximate equilibrium converge to an equilibrium in the original game[6].

1. Start with the initial guess of a probability that an offer z occurs in equilibrium if the state $s = s_1$: $\psi_0 : \mathbf{Z} \to [0,1]$, $\sum_{z \in \mathbf{Z}} \psi_0(z) = 1$, and $\psi_0(z) > 0$ for all z.

[6] See Section 3 of Fudenberg-Levine (1995) for a formal statement and a proof.

2. For any offer $z = (z^1, z^2)$ use Bayes' rule to find a posterior belief of any agent with a prior belief δ who receives an offer z:

$$\delta'(\delta, z) = \frac{\delta\psi_0((z^1, z^2))}{\delta\psi_0((z^1, z^2)) + (1-\delta)\psi_0((z^2, z^1))}.$$

If δ' falls outside of the grid point, we round it to the closest point on Δ. Since $\psi_0(z) > 0$ for all z, this rule is well defined.

3. Find the probability χ that an offer z is accepted in state 1. $\chi : \mathbf{Z} \to [0,1]$; $\chi(z) = \sum \Gamma(x, \delta | s_1)$ where the summation is over all $(x, \delta) \in \mathbf{X} \times \Delta$ s.t. $W(x + z, \delta'(\delta, z)) \geq W(x, \delta'(\delta, z))$.

4. Use Bayes' rule to find a posterior of the agent who makes the offer z if such an offer is accepted, δ_a, and a posterior if it is rejected, δ_r:

$$\delta_a(\delta, z) = \begin{cases} \frac{\delta\chi((z^1,z^2))}{\delta\chi((z^1,z^2))+(1-\delta)\chi((z^2,z^1))}, & \text{if } \delta\chi((z^1, z^2)) + (1-\delta)\chi((z^2, z^1)) > 0 \\ \delta, & \text{otherwise} \end{cases}$$

$$\delta_r(\delta, z) = \begin{cases} \frac{\delta(1-\chi((z^1,z^2)))}{\delta(1-\chi((z^1,z^2)))+(1-\delta)(1-\chi((z^2,z^1)))}, \\ \text{if } \delta(1 - \chi((z^1, z^2))) + (1-\delta)(1 - \chi((z^2, z^1))) > 0 \\ \delta, \text{ otherwise.} \end{cases}$$

If δ'' falls outside of the grid point, we round it to the closest point on Δ.

5. Find a utility $w(z; x, \delta)$ of the agent (x, δ) if he makes an offer z:

$$\begin{aligned} w(z; x, \delta) &= (\delta\chi((z^1, z^2)) + (1-\delta)\chi((z^2, z^1)))W(x - z, \delta_a(\delta, z)) \\ &\quad + (1 - (\delta\chi((z^1, z^2)) + (1-\delta)\chi((z^2, z^1))))W(x, \delta_r(\delta, z)) \end{aligned}$$

If the offer $(x - z) \notin \mathbf{X}$, let $w(z; x, \delta)$ be a large negative number, $-\underline{w}$.

6. Define a strategy of an agent with (x, δ) as $\sigma_m(z; x, \delta)$:

$$\sigma_m(z; x, \delta) = \frac{\exp(\kappa w(z; x, \delta))}{\sum_{z' \in \mathbf{Z}} \exp(\kappa w(z'; x, \delta))} \tag{32}$$

7. Find a probability of each offer $\sigma_m(z) = \sum_{(x,\delta) \in \mathbf{X} \times \Delta} \sigma_m(z; x, \delta)$. If $||\sigma_m - \psi_0||$ is less than the chosen precision, finish the procedure. Otherwise, let $\psi_1 = \frac{1}{2}\psi_0 + \frac{1}{2}\sigma_m$ and go to Step 1 (for subsequent iterations use $\psi_{n+1} = \frac{n}{n+1}\psi_n + \frac{1}{n+1}\sigma_m$ and repeat the procedure until $||\sigma_m - \psi_n||$ is less than the chosen precision).

In the procedure above, (32) ensures that, for all z, $\sigma_m(z) > 0$ and, since $\psi_0(z) > 0$, $\psi_n(z) > 0$ for all z, n. This ensures that Bayes rule for updating agent's beliefs in Step 2 is well defined.

In computations in Section **??** we further reduce computational complexity by restricting out of the equilibrium beliefs for some offers. We start by considering what is the *lowest* probability that an offer z can be accepted in *any* equilibrium. This probability, $\chi^{\min}(z)$ is defined as $\chi^{\min}(z) = \sum \Gamma(x, \delta|s_1)$ where summation is over all $(x, \delta) \in \mathbf{X} \times \Delta$, s.t. $\min_{\tilde{\delta} \in [0,1]} \left\{ W(x+z, \tilde{\delta}) - W(x, \tilde{\delta}) \right\} > 0$. Next, we follow Steps 3-5 to compute $w(z; x, \delta)$. We define $\sigma_m(z; x, \delta) = 1$ if $z = \arg\max_{z'} w(z'; x, \delta)$ and 0 otherwise and set $\chi_0(z) = \sum_{(x,\delta) \in \mathbf{X} \times \Delta} \sigma_m(z; x, \delta)$. Then we restrict the set of allowed offers to $\tilde{\mathbf{Z}} \equiv \{z \in \mathbf{Z} : \chi_0(z) > 0\}$. With these restrictions we use the iterative procedure described above. This procedure restricts all out of equilirbium beliefs to $\arg\min_{\delta'} \left\{ W(x+z, \delta') - W(x, \tilde{\delta}) \right\}$. Any offer in a set $\tilde{\mathbf{Z}}$ is accepted at least with a probability $\chi^{\min}$, which means that any offers in a set $\mathbf{Z} \backslash \tilde{\mathbf{Z}}$ are dominated by some offer in a set $\tilde{\mathbf{Z}}$ both on and off the equilibrium path

9.2 Finding a sequence of payoffs $\{W(\cdot, \cdot, t)\}_{t=1}^{\infty}$ and an equilibrium of the dynamic game

To compute an equilibrium of a dynamic game, we truncate the game at period T. We assume that if the game has not ended before period T, it ends with probability 1 in period $T+1$.

1. Make a guess on the distribution of beliefs and endowments $\left\{ \Gamma_t^0(\cdot, \cdot|s_1) \right\}_{t=1}^{T}$.

2. Let $W_{T+1}^0(\cdot, \cdot) = U(\cdot, \cdot)$. Use the procedure in Section 9.1 to compute equilibrium strategies for a static game with a payoff W_{T+1}^0 and distribution Γ_T^0. Obtains functions ψ, σ_m, $\chi^{\min}$, w.

3. Compute the payoff at the beginning of the period T. For this purpose, let W^m and W^r be, resectively, the payoffs the agents who make and and receive offers. Then

$$W^r(x, \delta) = \sum_{z \in \tilde{\mathbf{Z}}} \psi(z) \max \left\{ W_{T+1}^0 \left(x+z, \delta'(\delta, z) \right), W_{T+1}^0 \left(x, \delta'(\delta, z) \right) \right\}$$

   For any $\Gamma(x, \delta|s_1) > 0$ compute utility of the agent who makes an offer as

$$W^m(x, \delta) = \sum_{z \in \tilde{\mathbf{Z}}} \sigma_m(z; x, \delta) w(z; x, \delta)$$

   or,

$$\begin{aligned} W^m(x, \delta) \quad = \quad & \max\{ \max_{z \in \mathbf{Z} \backslash \tilde{\mathbf{Z}}} (\delta \chi^{\min}((z^1, z^2)) + (1-\delta)\chi^{\min}((z^2, z^1))) W(x-z, \delta_a(\delta, z)) \\ & + (1 - (\delta \chi^{\min}((z^1, z^2)) + \\ & + (1-\delta)\chi^{\min}((z^2, z^1)))) W(x, \delta_r(\delta, z)), \max_{z \in \tilde{Z}} w(z; x, \delta)\} \end{aligned}$$

   The beginning of period T payoff is then $\frac{1}{2} W^m + \frac{1}{2} W^r$.

4. Set $W_T^0 = \gamma \left(\frac{1}{2} W^m + \frac{1}{2} W^r \right) + (1-\gamma) U$, and returnt to Step 2 until the whole sequence $\left\{ W_t^0 \right\}_{t=1}^{T}$ is computed.

5. Start with the initial distribution $\Gamma_1(\cdot,\cdot|s_1)$ and W_1^0 from Step 2 and compute the equilibrium in a one shot game using the algorithm in Section 9.1. Compute

$$\begin{aligned}
\Gamma_2^1(\tilde{x},\tilde{\delta}|s_1) &= \frac{1}{2}\sum_{\substack{\{x,\delta,z:x-z=\tilde{x}\\ \delta_a(\delta,z)=\tilde{\delta}\}}} \sigma^m(z;x,\delta)(\delta\chi((z^1,z^2))+(1-\delta)\chi((z^2,z^1)))\Gamma_1(x,\delta|s_1) \\
&+\frac{1}{2}\sum_{\substack{\{x,\delta,z:x=\tilde{x}\\ \delta_r(\delta,z)=\tilde{\delta}\}}} \sigma^m(z;x,\delta)(1-(\delta\chi((z^1,z^2))+(1-\delta)\chi((z^2,z^1))))\Gamma_1(x,\delta|s_1) \\
&+\frac{1}{2}\sum_{\substack{\{x,\delta,z:\delta'(\delta,x)=\tilde{\delta},x+z=\tilde{x}\\ W(x+z,\delta'(\delta,z))\geq W(x,\delta'(\delta,z))\}}} \chi(z)\Gamma_1(x,\delta|s_1) \\
&+\frac{1}{2}\sum_{\substack{\{x,\delta,z:\delta'(\delta,x)=\tilde{\delta},x=\tilde{x}\\ W(x+z,\delta'(\delta,z))<W(x,\delta'(\delta,z))\}}} \chi(z)\Gamma_1(x,\delta|s_1)
\end{aligned}$$

Here, the first term is the transition probabilities of all makers whose offers are accepted, the second term is the transition probabilities of all makers whose offers are rejected, the third term is transition probabilities of all receivers who accept offers and the fourth term is the transition probabilities of all receivers who reject offers. $\Gamma(x,\delta|s_2)$ can be obtained from $\Gamma(x,\delta|s_1)$ using symmetry of equilibrium.

6. Go to Step 5 until the whole sequence $\{\Gamma_t^1\}_{t=1}^{\infty}$ is computed.

7. If $||\Gamma^1-\Gamma^0||$ ($||\Gamma^{n+1}-\Gamma^n||$ in subsequent iterations) is less than chosen precision, finish the procedure. Otherwise, proceed to Step 1.

9.3 Further simplifications with exponential utility function

The procedure described above can be further simplified by assuming exponential utility function $u(x)=-\exp(-x)$ and allowing agents to have any (both positive or negative) x in all periods. In this case the strategies of any agent depend on (x^1-x^2,δ,t), which reduces the number of state variables. To see that this is the case, consider a payoff for any agent (x,δ) in period t by following some strategy σ:

$$\begin{aligned}
W(x,\delta,t)(\sigma) &= E\left\{\sum_{k=0}^{\infty}(1-\gamma)^k\left[\begin{array}{c}\pi(\delta_{t+k}(\sigma_{t+k}))u(x_{t+k}^1(\sigma_{t+k})) \\ +(1-\pi(\delta_{t+k}(\sigma_{t+k})))u(x_{t+k}^2(\sigma_{t+k}))\end{array}\right]|\Pr(s=s_1)=\delta\right\} \\
&= E\left\{\sum_{k=0}^{\infty}(1-\gamma)^k\left[\begin{array}{c}\pi(\delta_{t+k}(\sigma_{t+k}))u(x^1+\sum_{m=0}^k z_{t+m}^1(\sigma_{t+m})) \\ +(1-\pi(\delta_{t+k}(\sigma_{t+k})))u(x^2+\sum_{m=0}^k z_{t+m}^2(\sigma_{t+m}))\end{array}\right]|\Pr(s=s_1)=\delta\right\} \\
&= \exp(-x^2)E\left\{\sum_{k=0}^{\infty}(1-\gamma)^k\left[\begin{array}{c}\pi(\delta_{t+k}(\sigma_{t+k}))u((x^1-x^2)+\sum_{m=0}^k z_{t+m}^1(\sigma_{t+m})) \\ +(1-\pi(\delta_{t+k}(\sigma_{t+k})))u(\sum_{m=0}^k z_{t+m}^2(\sigma_{t+m}))\end{array}\right]|\Pr(s=s_1)=\delta\right\}
\end{aligned}$$

Consider any two strategies, σ' and σ'', s.t. $W(x,\delta,t)(\sigma')\geq W(x,\delta,t)(\sigma'')$ for some (x,δ). Since we do not impose bounds on asset holdings x_t, the same strategies σ' and σ'' are feasible for all agents. But then the last expression implies that $W(\tilde{x}, \delta, t)(\sigma') \geq W(\tilde{x}, \delta, t)(\sigma'')$ for all $\tilde{x}$ s.t. $\tilde{x}^1 - \tilde{x}^2 = x^1 - x^2$.

References

[1] Amador, Manuel and Pierre Olivier Weill (2007). "Learning from Private and Public Observation of Other's Actions. Unpublished.

[2] Amador, Manuel and Pierre Olivier Weill (2008). "Learning from Prices: Public Communication and Welfare". Unpublished.

[3] Blouin, Mar R., and Roberto Serrano (2001), "A Decentralized Market with Common Values Uncertainty: Non-Steady States," Review of Economic Studies 68: 323–346.

[4] Brunnermeier, Markus (2001). Asset Pricing under Asymmetric Information. Oxford University Press.

[5] Dubey, Pradeep, John Geanakoplos and Martin Shubik (1987), "The Revelation of Information in Strategic Market Games: A Critique of Rational Expectation Equilibrium," Journal of Mathematical Economics 16: 105–137.

[6] Duffie, Darrell, Nicolae Garleanu and Lasse Heje Pedersen (2005), "Over-the-Counter Markets," Econometrica 73: 1815–1847.

[7] Duffie, Darrell, Nicolae Garleanu and Lasse Heje Pedersen (2006). "Valuation in Over-the-Counter Markets". Forthcoming in Review of Financial Studies.

[8] Duffie, Darrell, Gaston Giroux and Gustavo Manso (2007), "Information Percolation," working paper.

[9] Duffie, Darrell, and Gustavo Manso (2007), "Information Percolation in Large Markets," American Economic Review Papers and Proceedings 97: 203–209.

[10] Gale, Douglas (1987), "Limit Theorems for Markets with Sequential Bargaining," Journal of Economic Theory 43: 20–54.

[11] Glosten, Lawrence R., and Paul R. Milgrom (1985), "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Infromed Traders," Journal of Financial Economics 14: 71–100.

[12] Grossman, Sanford, and Joseph Stiglitz (1980), "On the Impossibility of Informationally Efficient Markets," American Economic Review 70: 392–408.

[13] Lagos, Ricardo (2007), "Asset Prices and Liquidity in an Exchange Economy", working paper.

[14] Lagos, Ricardo and Guillaume Rocheteau (2007), "Liquidity in Asset Markets with Search Frictions", working paper.

[15] Lagos, Ricardo, and Guillaume Rocheteau and Pierre Olivier Weill (2007), "Crashes and Recoveries in Illiquid Markets", working paper.

[16] Milgrom, Paul R., and Nancy Stokey (1982), "Information, Trade and Common Knowledge," Journal of Economic Theory 26: 17–27.

[17] Rogers, L. C. G and David Williams (2001). Diffusions, Markov Processes and Martingales: Volume 2, Itô Calculus (Cambridge Mathematical Library). Cambridge University Press.

[18] Wolinsky, Asher (1990), "Information Revelation in a Market with Pairwise Meetings," Econometrica 58: 1–23.

[19] Vayanos, Dimitri (1998). Transaction Costs and Asset Prices: A Dynamic Equilibrium Model. Review of Financial Studies 11, 1–58.

[20] Vayanos, Dimitri and Pierre Olivier Weill (2007). "A Search-Based Theory of the On-the-Run Phenomenon". Forthcoming, Journal of Finance.

[21] Weill, Pierre Olivier (2007). "Liquidity Premia in Dynamic Bargaining Markets". Forthcoming, Journal of Economic Theory.