



HAL
open science

Learning by Trading in Infinite Horizon Strategic Market Games with Default

Sonja Brangewitz, Gaël Giraud

► **To cite this version:**

Sonja Brangewitz, Gaël Giraud. Learning by Trading in Infinite Horizon Strategic Market Games with Default. 2012. halshs-00747899

HAL Id: halshs-00747899

<https://shs.hal.science/halshs-00747899v1>

Submitted on 2 Nov 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.




Documents de Travail du Centre d'Économie de la Sorbonne

C
E
S

W
o
r
k
i
n
g

P
a
p
e
r
s



Learning by Trading in Infinite Horizon Strategic Market Games with Default

Sonja BRANGEWITZ, Gaël GIRAUD

2012.62



Learning by Trading in Infinite Horizon Strategic Market Games with Default*

Sonja Brangewitz[†] and Gaël Giraud[‡]

September 20, 2012

Abstract

We study the consequences of dropping the perfect competition assumption in a standard infinite horizon model with infinitely-lived traders and real collateralized assets, together with one additional ingredient: information among players is asymmetric and monitoring is incomplete. The key insight is that trading assets is not only a way to hedge oneself against uncertainty and to smooth consumption across time: It also enables learning information. Conversely, defaulting now becomes strategic: Certain players may manipulate prices so as to provoke a default in order to prevent their opponents from learning. We focus on learning equilibria, at the end of which no player has incorrect beliefs — not because those players with heterogeneous beliefs were eliminated from the market (although default is possible at equilibrium) but because they have taken time to update their prior belief. We prove a partial Folk theorem à la Wiseman (2011) of the following form: For any function that maps each state of the world to a sequence of feasible and strongly individually rational allocations, and for any degree of precision, there is a perfect Bayesian equilibrium in which patient players learn the realized state with this degree of precision and achieve a payoff close to the one specified for each state.

Keywords and Phrases: Strategic Market Games, Infinite Horizon, Incomplete Markets, Collateral, Incomplete Information, Learning, Adverse Selection

JEL Classification Numbers: C72, D43, D52, G12, G14, G18

*We benefited from comments from the participants of the Workshop on Markets and Games at the Research Center for Games and Economic Behavior (RCGEB), Shandong University, (China), October 2011, 11th SAET Conference, 2011 (Portugal), the Workshop on Strategic Interactions and General Equilibrium: Theories and Applications II, 2011 at Université Paris Ouest Nanterre (France). Discussions with Jean-Marc Bonnisseau, Claus-Jochen Haake, Leonidas C. Koutsougeras, Enrico Minelli and Walter Trockel were especially illuminating. All remaining errors are ours.

[†]Department of Economics, University of Paderborn, Warburger Str. 100, 33098 Paderborn, Germany; sonja.brangewitz@wiwi.upb.de. An earlier version of this article has been part of the dissertation Brangewitz (2012). This research was mainly carried out at the Institute of Mathematical Economics, Bielefeld University, and the Centre d'Economie de la Sorbonne, Université Paris 1 Panthéon Sorbonne, while being a member of the International Research Training Group “Economic Behavior and Interaction Model” (EBIM) financed by the German Research Foundation (DFG) under contract GRK 1134/2. Financial support through the Franco-German University (DFH – UFA) is gratefully acknowledged. This work was partially supported by the German Research Foundation (DFG) within the Collaborative Research Centre “On-The-Fly Computing” (SFB 901).

[‡]CNRS, CES, Paris School of Economics, Paris, France; ggiraud@parisschoolofeconomics.eu.

1 Introduction

We explore the consequences of dropping the perfect competition assumption in an infinite horizon model populated by infinitely lived traders. We study a strategic market game with an infinite horizon and collateralized assets. Regarding the information structure, we consider a middle ground between the two strands of literature cited above: Players start the game with the same prior, but along the play they gather different information so that in general, information will be asymmetric among them. Moreover, although they may have heterogeneous beliefs during a play, they cannot be simply driven out of the market thanks to the collateral constraints, as in the model of Cao (2011). We rather focus on equilibria where players strategically learn the true state of the world. As a consequence, in general, the degree of belief heterogeneity will be varying over time, as in He and Xiong (2010). In Cao (2011), there is no learning but, as noticed by the author (page 9), learning could be added to his framework by considering the current beliefs of agents as additional state variables. As in Blume and Easley (2006), agents who learn slowly would disappear under complete markets. However they would all survive under collateral constraints. The dynamics of asset prices described in Cao (2011) are then interpreted as the short-run behavior of asset prices in the economy with learning. Here, we make learning explicit, and, above all, strategic — so that even the speed of learning is now endogenous. This enables us to deal with the long-run behavior of asset prices and trades under heterogeneous beliefs.

An infinite market game with dependent Markovian types

The model considered here extends the finite horizon case *without* default considered in Giraud and Weyers (2004) and the finite horizon *with* default examined in Brangewitz (2012, Chapter 4). In both papers, the uncertainty is only on future endowments, while here we allow for uncertainty on endowments, utilities and asset returns.¹ Moreover, the authors restricted themselves to a very specific game-theoretic setup: One with partial monitoring (players condition their actions on the public history of prices but not on traded quantities, and on the private history of their own individual trades) and ex ante evaluation of each player's payoff — that is, when contemplating a counterfactual, a player considered only the ex ante impact of her deviation with respect to the expectation operator computed thanks to some prior belief over the whole event tree. Everything being computed ex ante, there was no learning process during the play of the game.²

The informational setup considered in this paper is the following. The probability distribution according to which uncertainty realizes in each period is a (stationary) Markov chain. This is much more general, of course, than the iid framework initially considered in the literature devoted to repeated games of adverse selection (Fudenberg et al., 1994). As emphasized, e.g., by Escobar and Toikka (2012), the serial dependence between the states of nature chosen in each period implies that players have private information about the distribution of future payoffs. More generally, any information disclosed at time t provides valuable information for latter stages, as in Golosov et al. (2009). This leads to the possibility of signaling: A player may now be tempted to alter her behavior to influence the other players' beliefs about her own type. Yet, whenever the Markov chain

¹See Thomas (1995) for an example of a general equilibrium model where uncertainty affects consumer's future utilities.

²Here, as in the literature just cited, we confine players to pure strategies. Adding mixed strategies would not necessarily impair our result but would introduce subtle additional considerations similar to those analyzed in De Meyer and Saley (2003). There, a market game with two players and lack of information on one side is studied. It is proven that if mixed strategies are allowed, it may be in the interest of the informed player to play randomly so as to hide her information (such ideas clearly go back at least to Hart, 1985). De Meyer and Saley (2003) show that such a behavior may induce a price volatility capable of providing a theoretical ground to the Brownian motion assumption common to a vast part of the literature in finance.

is irreducible, this information becomes eventually valueless since the time-average of any player's signal should reflect the invariant measures of the states of nature, or else it will be disregarded anyway.³ At variance with Escobar and Toikka (2012) (and with most of the burgeoning literature devoted to self-correlated types in repeated games), however, we drop the assumption that the evolution of individual types is independent across players.⁴ More specifically, here the transition matrix itself is chosen at random once and for all at the start of play, and the investors do not observe the random choice of this *state of the world*. The distribution of initial endowments and Bernoulli utility functions are chosen randomly in each period according to the state of the world, and privately observed. Neither do they observe any other player's action — so that markets keep with the anonymity of trades. Rather, the players have a common prior⁵ over the finite set of possible Markov chains (states of the world), and they have various ways of updating their belief over time. First, each player observes her own initial endowment and utility function in each period — both are realizations of random variables whose distribution depends on the state of nature, hence on the state of the world. Furthermore, each player observes the return of each financial asset she owns in her portfolio (either as a creditor or a debtor) unless this asset defaults on its promise. In the latter case, the collateral is forfeit but the precise delivery of the return remains unknown. For simplicity, no utility punishment is imposed on defaulting debtors.

For investors to be able to learn the state, we flesh out the general equilibrium skeleton with a strategic market game.⁶ Collateral requirements for financial assets are introduced as in Geanakoplos and Zame (2007) and the related literature. Nevertheless, players can manipulate their opponents' information by influencing publicly announced prices. Despite the risk of information manipulation, however, those traders with incorrect beliefs can realize their mistake along the play of the game, and strategically learn the state of the world. We therefore focus on learning equilibria, at the end of which no player has incorrect beliefs — not because they were eliminated from the market (although default *is* possible at equilibrium) but because they have taken time to cleverly update their prior belief. Our main result is a partial Folk theorem à la Wiseman (2011): For any function that maps each state of the world to a sequence of feasible and strongly individually rational allocations (precise definitions are given in section 3), and for any degree of precision, there is a perfect Bayesian equilibrium in which patient players learn the realized state with this degree of precision and achieve a payoff close to the one specified for each state. Hence, within this class of equilibria, no player with incorrect beliefs stays on the market in the long run, provided she is patient enough — thus confirming Friedman's (1953) hypothesis but with a completely different argument. Obviously, the infinite repetition of the autarkic Nash equilibrium, is a perfect Bayesian equilibrium. This shows that our result is only partial: There are plenty other equilibria where no learning occurs.

The double role of financial assets

Along the path of equilibria considered in this paper, players with incorrect beliefs can now learn the state of the world (hence better forecast their future payoffs) through coordinated experimentation by trying different action profiles, observing the resulting payoff realiza-

³See, e.g., RSV (2010).

⁴Two other differences with Escobar and Toikka (2012) are the following: We do not permit cheap talk and we drop the perfect monitoring restriction. That is, here players have no other way to communicate but through market prices. Last, we prove an approximate Folk theorem as the discount factor goes to one while Escobar and Toikka (2012) only prove that Pareto-efficient streams of payoffs can be approximated by Bayesian perfect equilibria.

⁵See the footnote on page 15 for an argument in favor of this assumption. This does not mean, of course, that players start with symmetric information.

⁶See Giraud (2003) for an introduction.

tions and updating their beliefs about the region of the event tree where they are currently located. Financial assets therefore play a double role: On the one hand they serve as means for reallocating one's resources in face of risky future events, on the other hand they can have a function analogous to that of "arms" in the multiarmed bandit problem (Rothschild, 1974). Each player's prior beliefs about the return distribution induce subjective payoff expectations for each asset, but the asset with the highest subjective expected payoff may not be the best one to choose: A trader may prefer to sacrifice expected returns in the short run to gain some information that will prove useful in the long run. Since there are several traders meeting on the same market, however, the situation becomes more complicated: Experimentation has to be somehow coordinated to be effective, since each trader must deliver information through a specified action and strategic considerations may interfere with learning.

As an example, suppose that two traders must decide repeatedly whether or not to exchange some given financial asset. In each period, the buyer incurs a cost π (the security's price), but the next period, the seller incurs the risk of having to pay a return $a > 0$ ("bad state") to the buyer, or to receive $b > 0$ from her ("good state" from the seller's viewpoint). It is worthwhile for the players to trade only if the discounted mean value of the payoff is greater than π for the buyer and the mean value of losses is smaller than π for the seller. But the only way to find out the mean value is to experiment by effectively trading in order to learn across time what the next return of this very asset will be.

Consequently, default becomes a way to prevent one's opponents from learning. Recall, indeed, that when entering markets in a given node of the event tree, players do not know the current state of nature. Therefore, although each of them knows the contract governing each financial asset, and despite the fact that current prices are publicly announced, the actual payoff of any given asset j remains unknown to any trader i , unless she effectively trades this very asset. By observing the actual final payoff of the assets she trades, player i will be able to revise her beliefs. This learning process, however, strikes a major rock: If asset j defaults,⁷ none of these two contractants will ever be able to observe the actual payoff of asset j . All they know is that this payoff should have been at least as large as what was effectively paid. As a consequence, a trader may prefer to sacrifice expected returns in the short run by enduring default in order to *prevent her opponents* from gaining some information that will prove useful to them in the long run. And in the same way as learning by trading is complexified by the very fact that several traders are competing against each other, strategic default is made more complicated by the fact that prices (whose levels will imply whether default occurs or not) are co-determined simultaneously by all the traders.

The piece of good news provided here is that as long as it is compatible with our key Informativeness Assumption (IA, to be described in section 4 below), the possibility of default does not prevent investors from learning the state. To put it differently, the Shapley-Shubik trading mechanism provides enough structure⁸ in the game so that despite price manipulation, infinite-horizon incomplete markets equilibria may be fully revealing in the following sense: A broad set of allocations can be approximated by Bayesian perfect equilibria, provided players are sufficiently patient. This is a strategic counterpart to the general equilibrium literature with real assets (but no default), where generically every equilibrium is fully revealing (Radner 1979, Duffie and Shafer 1985). In addition to the deep difference between our imperfectly competitive approach and the perfect competition

⁷Observe (see equation (D) *infra*) that the condition under which the seller of an asset defaults depend only on anonymous prices. As a consequence, if one seller of asset j defaults, so do all its sellers. In this sense, we can speak of "the default of an asset".

⁸The most obvious one is that the analog of the "minimax" strategy (which consists here in preventing a deviator from purchasing more than $\varepsilon > 0$ of each good) can be made independent from the identity of the deviator. This key property will free us from the need to use the tools of Blackwell approachability (B56 (1956)).

hypothesis, a careful reading of our proof, however, shows that our result goes through in the nominal asset case as well.⁹ Therefore, from the strategic point of view adopted in this paper, there is no essential difference between real and nominal assets. This contrasts with the negative results obtained in the perfectly competitive general equilibrium literature with incomplete markets of nominal assets (see Rahi 1995 and the references therein).¹⁰

Last, perfect competition with infinite horizon and incomplete markets faces an important stumbling block for existence, due to the possibility of Ponzi schemes at equilibrium. As a consequence, the literature devoted to this setting usually relies on some transversal budget constraint in order to forbid such Ponzi schemes (see, e.g., Florenzano and Gourdel 1996). On the other hand, when collateral requirements are added, Araujo et al. (2002) show that no Ponzi scheme arises at equilibrium. In our imperfectly competitive setup, there is no need for any transversal budget constraint, even when markets are complete. Due to the finite number of investors, indeed, a Ponzi scheme would require at least one player to borrow money from at least one other player during an infinite number of periods. The lender would clearly better not lend her money so many times — hence, participating in a Ponzi scheme cannot be part of everyone's best reply (see, e.g., O'Connell and Zeldes 1988). This is true with and without collateral constraints.

The Informativeness Assumption (IA)

The kind of uncertainty under scrutiny in this paper affects each investor's initial endowments, her utility function, and the returns of financial assets. This setting captures many aspects extensively studied in the literature in terms of adverse selection. One key assumption in this paper (the Informativeness Assumption, (IA)) can be stated as follows: Observing the realization of one's (random) initial endowments, one-period (random) utility functions and (strategically determined) final allocations together with *all* the assets' returns, individual final allocations and (public) prices suffices for every single trader to learn the true state of the world in the long run with probability arbitrarily close to one. Needless to say, this assumption is far from being sufficient to guarantee *a priori* that every player will always learn the true state with arbitrary accuracy: For that purpose, she needs to be able to keep every asset in her portfolio in every period; she may be diverted by the strategic signaling of her opponents; the learning process must remain compatible with the equilibrium conditions, hence should not involve too deep losses. On the other hand, (IA) is verified in a number of important instances:

Arrow securities

(IA) is clearly satisfied when the asset structure is that of Arrow securities, where each security pays off one in one single state. In this case, observing assets' returns suffices to identify the Markov chains' realization after each round of trade (even without taking account of prices or of one's private knowledge gained by observing endowments and stage payoffs). After a sufficiently long time, if every trader succeeds in observing every asset's return, the true state of the world will become common knowledge. Notice, however, that even in this polar case, full revelation at a strategic equilibrium is not straightforward, and there is something to be proven: Indeed, our argument requires that *every* trader be able to trade *every* Arrow security in *every* period. If one of them fails to observe all the

⁹The proof is actually even simpler. This is one reason why we have treated the real asset case.

¹⁰As a side consideration, our approach may shed some light on the current debate about dark pools (see Zhu 2011). Dark pools are trading systems that do not display their orders to the public markets. A recent literature investigates whether dark pools harm price discovery. In light of our anonymous trading assumption, our result can be interpreted as showing that as long as only market orders are allowed, dark pools do not prevent intermediaries from correctly learning the state of the world. Further investigation in this direction would require to refine the market microstructure and to allow players to send limit-price (not just market) orders to the clearing house.

assets' returns in certain periods, then she might not draw the right conclusion about which Markov chain is driving uncertainty, so players cannot coordinate on any state-dependent equilibrium path. On the other hand if, say, only the riskless asset (delivering the same return in every state) is marketed, then observing assets' returns does not provide any information.

Akerlof's model

Akerlof's (1970) model of used cars is a static one. Its extension to our intertemporal framework can easily be interpreted as verifying (IA). Suppose, indeed, that the quality index, s , of a car is an integer belonging to $[1, 10]$. s is distributed according to the Markov chain ω . As quality of a car is indistinguishable beforehand by the buyer (due to the asymmetry of information), incentives exist for the seller to pass off low-quality goods as higher-quality ones. The buyer, however, takes this incentive into consideration and takes the quality of the goods to be uncertain. Only the average quality of the goods will be considered, which in a one-shot setup will have the side effect that goods that are above average in terms of quality will be driven out of the market. In our multi-period setting, however, this need not occur: Each time t , the seller receives a new (random) endowment of used cars. Each period, the buyers are informed *ex post* (through their stage payoff) about the actual quality, s , of the car they have bought. Across time, they may learn the transition matrix ω , hence anticipate the distribution of s in the future. Our main result then says that the observation of prices and private knowledge enables actors on the market for used cars to enforce a large set of effective trades. This sharply contrasts with Akerlof's conclusion that the market for used cars should collapse.

Moral hazard.

Since investors take privately observed actions affecting their initial endowments and portfolios, our paper is also linked to the literature on moral hazard. The differences in information and the signaling aspects of the present work are, for example, related to the job market signaling model of Spence (1973) or the competitive insurance market considered in Rothschild and Stiglitz (1976). However, we depart from the classical principal agent model in as much every individual may act as a seller or a buyer (or both simultaneously), and this on commodity as well as on asset markets. Therefore, we cannot impose, for example, that a seller always is less informed than a buyer or vice versa. Finally, we consider only finitely many players. Our setup therefore sharply differs from the perfectly competitive case studied in the seminal papers by Prescott and Townsend (1984a,b) or, more recently, by Acemoglu and Simsek (2010). In particular, we get a wide range of equilibria including allocation streams that are Pareto-optimal and others that are dominated. Thus, our result stands at distance both from the generic inefficiency obtained by Greenwald and Stiglitz (1984) or Arnott and Stiglitz (1986, 1990, 1991), and from the more positive results obtained by Acemoglu and Simsek (2010).

The paper is organized as follows: First, in Section 2, we describe the infinite horizon economy and its associated strategic market game. Section 3 focuses on a particularly important subclass of allocations that plays a key role in the sequel. The next section proves a (partial) Folk theorem with incomplete information. Section 5 draws a non-convergence result from the previous theorem. The last section presents the conclusion.

2 The Markov Strategic Market Game with Collateral

The underlying economy is a rather standard model with infinitely-lived investors (see Kubler and Schmedders (2003) or Cao (2011)).

2.1 The Markov Economy

The environment

Time is discrete. In each period, t , the state of nature in the next period, $t+1$, is chosen using a Markov transition matrix, ω , with a finite set of possible states of nature $\mathcal{S} = \{1, \dots, S\}$. As in Magill and Quinzii (1994) and the subsequent literature, time, uncertainty and the revelation of information are described by an event tree, i.e., a directed graph $(\mathbf{D}, \mathcal{A})$ consisting of a set \mathbf{D} of nodes and a set $\mathcal{A} \subset \mathbf{D} \times \mathbf{D}$ of (oriented) arcs. A node $z \in \mathbf{D}$ can be interpreted as a date-event pair (s_{t-1}, s) , where $t \geq 1$ is the minimal length of a walk between the root of the event tree, denoted by z_0 , and z , $s_{t-1} \in \prod_{t'=1}^{t-1} \mathcal{S}$ is the sequence of realizations of the state of nature up to $t-1$ and $s \in \mathcal{S}$ is the last state in t . Let $\tau(z)$ be the time at which node z is reached, i.e. $\tau : \mathbf{D} \rightarrow \mathbb{N}$ such that $z = (s_{t-1}, s) \mapsto t$. Define a partial order \geq on \mathbf{D} by $z = (s_{t-1}, s) \geq z' = (s_{t'-1}, s')$ if and only if there is a walk from z' to z . Of course, if $z \neq z'$ and $z \geq z'$, then $z > z'$. The unique predecessor of z is denoted by $z^- = (s_{t-2}, s')$.¹¹ The set of immediate successors of z , denoted by z^+ , is the set of nodes that are adjacent from z . For any node $z \in \mathbf{D}$, the set of all nodes with $z' \geq (>)z$ is denoted by $\mathbf{D}(z)$ ($\mathbf{D}(z)^+$) and is itself a tree with root z .

Recall that given some Markov chain ω , $\mu_\omega \in \Delta(\mathcal{S})$ is an *invariant measure* of ω if

$$\mu_\omega(s) = \sum_{s'} \omega_{s's} \mu_\omega(s') \quad \forall s \in \mathcal{S}.$$

We assume that the Markov chain ω is **irreducible** and **aperiodic**.¹² We furthermore assume throughout the paper that every state $s \in \mathcal{S}$ is positive recurrent.¹³ Therefore, ω admits a unique invariant measure.

A state of the *world* corresponds to the transition matrix, ω , of dimension $S \times S$ that is chosen once and for all at time 0 before the start of the play. We assume that there are finitely many states of the world, $\omega \in \Omega$.

Consumption goods and financial assets

Let us consider a pure exchange economy \mathcal{E} with a finite set, $\mathcal{N} = \{1, \dots, N\}$, of investors, L consumption goods, indexed by ℓ , and J short-term real assets, indexed by j . A financial asset $j \in \mathcal{J} := \{1, \dots, J\}$ is characterized by a tuple (z^j, A_j, C_j) consisting of three elements: an issuing node, promised deliveries and collateral requirements. The issuing node (a node in the tree \mathbf{D}) is denoted by z^j . We only consider short-term assets: The promised amount of goods is described by a function $A_j : \mathbf{D} \rightarrow \mathbb{R}_+^L$ such that $A_j(z) = 0$ for all $z \in \mathbf{D} \setminus (z^j)^+$ (i.e., for other nodes before the issuing node and two or more periods after the asset was issued, the promise is zero). For $z' \in (z^j)^+$, the promise $A_j(z')$ is the bundle of goods that a seller of asset j promises to deliver to a buyer of asset j in the next period following the issuing node z^j . The delivery, $p_z \cdot A_j(z)$, is assumed to be made in fiat money using spot prices, $p_z \in \mathbb{R}_+^L$. The vector $C_j \in \mathbb{R}_+^L$ is the amount of collateral needed at the issuing node, z^j , in order to back up the promised delivery A_j . Only consumption goods can serve as collateral.¹⁴ Commodities are assumed to be **storable** but not durable. Thus, they have to be consumed at the very date they enter the economy (as initial endowment), unless they

¹¹We define $s_{-1} = \emptyset$.

¹²A state $s \in \mathcal{S}$ has period k if any return to state s must occur in multiples of k steps. If $k = 1$, state s is said to be aperiodic. If every state $s \in \mathcal{S}$ is aperiodic, ω is said to be *aperiodic*. The Markov chain ω is *irreducible* if it is possible to connect every state $s \in \mathcal{S}$ with any other state $s' \in \mathcal{S}$ with positive probability.

¹³A state s is recurrent if, given that the chain starts in s , it will return to s in finite time with probability one. s is *positive recurrent* if, in addition, the expectation of this hitting time is finite.

¹⁴i.e., we do not introduce securities that are backed by other securities: Pyramiding is not allowed.

are stored as collateral. Tobacco and wine are examples of such commodities. Individuals are not allowed to consume a collateral. It is stored in a warehouse for one period. For simplicity, after having been stored for one period, a collateral must be consumed, otherwise it gets lost.¹⁵ For simplicity, at each node $z \in \mathbf{D}$ the same finite number of financial assets is issued.

2.2 The Strategic Market Game with Collateral

The players

Every player $i \in \mathcal{N}$ is characterized by a twice continuously differentiable, strictly increasing and concave per node utility function $u_z^i : \mathbb{R}_+^L \rightarrow \mathbb{R}$ and a non-zero initial endowment in consumption goods $w_z^i \in \mathbb{R}_+^L \setminus \{0\}$ at every node $z \in \mathbf{D}$. We assume that $(u_z^i(\cdot))_z$ are uniformly bounded below for all individuals i . Therefore, without loss of generality suppose $u_z^i(0) = 0$. Moreover, we assume that individual endowments are uniformly bounded above by some \bar{w} , across individuals and periods. Initial holdings of assets are 0.

We denote by $\mathcal{E}_z = \langle w_z^i, u_z^i(\cdot), (z^j, A_j, C_j)_{j|z^j=z} \rangle$ the finite-dimensional *one-shot economy* at node z .

The timing

The timing of events is as follows:

- (0) The state of the world $\omega \in \Omega$ is chosen (once and forever).
- (1) Individuals observe their initial endowments and their current (node-)utility function *privately*.
- (2) The strategic market game is played: Bids and offers on commodity and financial markets are submitted.
- (3) Resulting prices are *publicly* announced and final allocations are *privately* observed.
- (4) Beliefs about the state of the world are updated according to the information revealed from (1) to (3).
- (5) The next state of nature is chosen according to ω and the succeeding period starts at (1).

Given such a timing, let us illustrate the kind of informational structure faced by the players in our setup by means of a simple example.

Example 2.1 Suppose there are 2 players, 2 commodities and 3 states of nature with

- in state s_a , the individual characteristics of the two players are:

$$\begin{aligned} u_{s_a}^1(x, y) &= \sqrt{xy}, \quad \omega_{s_a}^1 = (\varepsilon, 1) \\ u_{s_a}^2(x, y) &= \sqrt{xy}, \quad \omega_{s_a}^2 = (1, \varepsilon) \end{aligned}$$

- in state s_b :

$$\begin{aligned} u_{s_b}^1(x, y) &= \sqrt{xy}, \quad \omega_{s_b}^1 = (\varepsilon, 1) \\ u_{s_b}^2(x, y) &= x + y, \quad \omega_{s_b}^2 = (1, 1) \end{aligned}$$

- in state s_c

¹⁵We could allow for a longer life expectancy of a collateral, say of length K , but at the cost of cumbersome notations. We thus take $K = 1$.

$$\begin{aligned} u_{s_c}^1(x, y) &= x + 2y, \quad \omega_{s_c}^1 = (2, 1) \\ u_{s_c}^2(x, y) &= x + y, \quad \omega_{s_c}^2 = (1, 1) \end{aligned}$$

with $0 < \varepsilon < 1$. In state s_a , player 2 knows the state by inspecting her own individual characteristics, but player 1 ignores whether $s = s_a$ or $s = s_b$. In state s_b , both players do not know the state of nature. In state s_c , player 1 knows the state but player 2 does not know whether $s = s_b$ or $s = s_c$. In state $s = s_a$ or $s = s_c$, the unique way for the uninformed player to refine her information partition about the state of nature consists in drawing some inference from the performed trades in commodities and assets and from the assets' returns. We therefore now turn to the description of the trading mechanism.

The trading mechanism

The trading microstructure is given by a market game à la Shapley-Shubik (see Shapley and Shubik (1977)). Each individual places for every consumption good $\ell \in \mathcal{L}$ at every node $z \in \mathbf{D}$ a bid $b_{z,\ell}^i$ and an offer $q_{z,\ell}^i$. The bid $b_{z,\ell}^i$ signals how much (in terms of fiat money) player i is willing to pay for the purchase of good ℓ , and the offer $q_{z,\ell}^i$ (in terms of physical commodities) is the amount she wants to sell. The price of good ℓ is given by:

$$p_{z,\ell} := \begin{cases} \frac{\sum_{i=1}^N b_{z,\ell}^i}{\sum_{i=1}^N q_{z,\ell}^i} & \text{if } \sum_{i=1}^N q_{z,\ell}^i > 0 \\ 0 & \text{otherwise.} \end{cases}$$

A trading post without trade is said to be closed.¹⁶

Similarly, at every node $z \in \mathbf{D}$ each player places a bid $\beta_{z,j}^i$ stipulating the amount of money she is ready to spend in buying asset j and offers for sale $\gamma_{z,j}^i$ units of this very asset. The asset's price is:

$$\pi_{z,j} := \begin{cases} \frac{\sum_{i=1}^N \beta_{z,j}^i}{\sum_{i=1}^N \gamma_{z,j}^i} & \text{if } \sum_{i=1}^N \gamma_{z,j}^i > 0 \\ 0 & \text{otherwise.} \end{cases}$$

When the promises are settled, a seller of the financial asset $j \in \mathcal{J}$ compares the value of the promise with the value of the collateral and pays back the minimal value:

$$D_{z',j} := \min \{ p_{z'} \cdot A_j(z'), p_{z'} \cdot C_j \} \quad (\text{D})$$

at node $z' \in (z^j)^+$. Hence, whether player i defaults or not is not the outcome of a deliberate decision of i alone, but depends upon the commodity price $p_{z'}$, which is strategically determined by all players' bids and offers posted at node $z' \in (z^j)^+$.

Feasible bids and offers

At every node $z \in \mathbf{D}$ and for every financial asset, player i needs to own the required amount of collateral, which depends on the quantity of the asset offered for sales and *not*

¹⁶By defining the price as zero when there is no offer on the market we follow Amir et al. (1990, p.128), Postlewaite and Schmeidler (1978, p.128), Peck et al. (1992, p.275) or Giraud and Weyers (2004, p.474), for example.

on the net trades.¹⁷ Assuming player i offers to sell $\gamma_{z,j}^i$ units of asset j at node z , then she needs to store $\gamma_{z,j}^i C_j \in \mathbb{R}_+^L$ as collateral.¹⁸

Feasible bids and offers must satisfy the following two constraints for all commodities ℓ :

$$\sum_{j=1}^J \gamma_{z,j}^i C_{j\ell} \leq w_{z,\ell}^i \quad (F1z)$$

and

$$q_{z,\ell}^i \leq \sum_{j=1}^J \gamma_{z-,j}^i C_{j\ell} + \Delta(F1z), \quad (F2z)$$

where $\Delta(F1z)$ stands for the difference between the right-hand side and the left-hand side of $(F1z)$. Inequality $(F1z)$ says that the collateral stored by i at node z must be taken out of her initial endowments. In particular, the collateral cannot consist of commodities that are already inherited from the past as collaterals. This formalizes our assumption that every commodity lives at most one period: Either it is consumed at the very period it enters into the economy (as initial endowment) or it is stored and consumed one period later. Notice that in the second period of a collateral's life, it may be traded by its owner and consumed by another player. Condition $(F2z)$ says that the offered amount of goods plus the amount of goods that must be stored as a collateral cannot exceed the initial endowment of player i at node $z \in \mathbf{D}$ plus the collateral that was put aside in the previous period. Of course, we impose:

$$q_{z,\ell}^i, b_{z,\ell}^i, \beta_{z,j}^i, \gamma_{z,j}^i \geq 0 \quad (F3z)$$

for all $\ell \in \mathcal{L}$, $j \in \mathcal{J}$.

The budget constraint

Player i also faces the following budget constraint on fiat money when placing bids and offers:

$$\begin{aligned} & \sum_{\ell=1}^L b_{\zeta,\ell}^i + \sum_{j=1}^J \beta_{\zeta,j}^i \\ & \leq \sum_{\ell=1}^L p_{\zeta,\ell} q_{\zeta,\ell}^i + \sum_{j=1}^J \pi_{\zeta,j} \gamma_{\zeta,j}^i + \sum_{j=1}^J \left(\theta_{\zeta-,j}^i - \varphi_{\zeta-,j}^i \right) D_{\zeta,j} \end{aligned} \quad (*_z^i 1)$$

for all $\zeta \leq z$ where $\theta_{\zeta-,j}^i$ denotes the final asset purchases and $\varphi_{\zeta-,j}^i$ the asset sales at node ζ^- (as will be defined below). Thus, by condition $(*_z^i 1)$ the total value of bids cannot exceed the amount of money player i can get given her sales and given the dividends received from her portfolio, $\theta_{\zeta-,j}^i - \varphi_{\zeta-,j}^i$. If $(*_z^i 1)$ is violated, say at node z , individual i is removed from the game for all subsequent nodes $\mathbf{D}^+(z)$, and all her goods are confiscated forever.¹⁹

¹⁷As discussed in Dubey and Geanakoplos (2003), netting before imposing the collateral requirement would suppress any constraint on the size of short sales. While this impairs the very existence of a competitive equilibrium (cf. Hart (1975)), it would make the proof of our partial Folk theorem only easier.

¹⁸Later, on page 11 when defining the final allocation in consumption goods, the collateral requirement is taken using the final asset sales, denoted by $\varphi_{z,j}^i$ and not directly on the offers $\gamma_{z,j}^i$.

¹⁹Therefore, for simplicity, we do not consider refinements of the solvency constraints such as the not-too-tight constraints aiming, among other things, at capturing reputation effects (see, e.g., ?).

The following condition will also be helpful: For every i :

$$\text{Either } \sum_{k \neq i} \gamma_{z,j}^k \neq 0 \quad \text{or} \quad \sum_{k \neq i} \beta_{z,j}^k \neq 0, \quad (*_z^i 2)$$

which says that there is at least one other individual on the bidding or on the offering side of the financial markets to trade with i .

Final allocations

After trading took place, player i 's holdings of asset $j \in \mathcal{J}$ are given by her sales

$$\varphi_{z,j}^i = \begin{cases} \gamma_{z,j}^i & \text{if } (*_z^i 1) \text{ and } (*_z^i 2) \text{ holds} \\ 0 & \text{otherwise} \end{cases}$$

and her purchases

$$\theta_{z,j}^i = \begin{cases} \frac{\beta_{z,j}^i}{\pi_{z,j}} & \text{if } (*_z^i 1) \text{ and } (*_z^i 2) \text{ hold and } \pi_{z,j} > 0 \\ 0 & \text{otherwise.} \end{cases}$$

Moreover, player i 's allocation of good $\ell \in \mathcal{L}$ available for consumption at the end of the current period at node z , is

$$x_{z,\ell}^i = \begin{cases} w_{z,\ell}^i + \sum_{j=1}^J \varphi_{z^-,j}^i C_{j\ell} - q_{z,\ell}^i + \frac{b_{z,\ell}^i}{p_{z,\ell}} - \sum_{j=1}^J \varphi_{z,j}^i C_{j\ell} & \text{if } (*_z^i 1) \text{ holds and } p_{z,\ell} > 0 \\ w_{z,\ell}^i + \sum_{j=1}^J \varphi_{z^-,j}^i C_{j\ell} - q_{z,\ell}^i - \sum_{j=1}^J \varphi_{z,j}^i C_{j\ell} & \text{if } (*_z^i 1) \text{ holds and } p_{z,\ell} = 0 \\ 0 & \text{otherwise.} \end{cases}$$

Remark 2.1 Condition $(*_z^i 2)$ emerges naturally once collateral requirements are introduced. Suppose, indeed, that individual i is the only one who wants to trade on the financial markets, i.e., $\sum_{k \neq i} \gamma_{z,j}^k = \sum_{k \neq i} \beta_{z,j}^k = 0$. Absent condition $(*_z^i 2)$, this individual could open the markets by bidding and offering strictly positive amounts of assets. By doing so, every player could store some collateral until the next period just by trading “with herself” today. We therefore require the collateral condition to be in force only when there is effective trade between player i and any other trader.

3 Interim Individual Rationality

3.1 Allowable strategies

The action set of player i at node z consists in feasible bids and offers:

$$\mathcal{A}_z^i = \left\{ (q_{z,\ell}^i, b_{z,\ell}^i)_{\ell \in \mathcal{L}}, (\gamma_{z,j}^i, \beta_{z,j}^i)_{j \in \mathcal{J}} \in \mathbb{R}_+^{2L} \times \mathbb{R}_+^{2J} \mid (F1z), (F2z) \text{ and } (F3z) \text{ are satisfied} \right\}.$$

Notice that \mathcal{A}_z^i depends upon z but not upon ω . Let $\mathcal{A}_z := \prod_{i=1}^N \mathcal{A}_z^i$.²⁰ The stage payoff of player i at node $z = (s_{t-1}, s)$ is given by the utility $u_z^i(x_z^i)$ she obtains from consumption.

Most importantly, and to keep the anonymity of markets, actions are *not* observed along the play of the game, which contrasts with the setting considered by Wiseman (2011), for

²⁰Note that the definition of an action set includes actions that possibly violate the budget constraint $(*_z^i 1)$ or $(*_z^i 2)$. An alternative would consist in incorporating these constraints into the very definition of a player's strategy set, which would lead to a generalized game as introduced by Debreu (1952) (see also Harker (1991) or Facchinei and Kanzow (2010)).

example. Prices are publicly observed by every player. At each node, every player also observes her own initial endowment, her utility function, her final allocation and also the returns of the assets present in her portfolio. Together with the public history of prices, these observations constitute the *private history* of player i . A *strategy* of player i consists in choosing an action at every node $z \in \mathbf{D}$ as a function of her own private history. Let H_z^i denote the set of possible private histories for individual i at node z , given by

$$H_z^i := \left\{ \left(w_{z'}^i, u_{z'}^i(\cdot), p_{z'}, \pi_{z'}, \varphi_{z'}^i, \theta_{z'}^i, \varphi_{z'}^i D_{z'}, \theta_{z'}^i D_{z'}, x_{z'}^i, w_z^i, u_z^i(\cdot) \right) \mid \forall z' < z \right\}.$$

The history at the root z_0 is given by $H_{z_0}^i = \{w_{z_0}^i, u_{z_0}^i(\cdot)\}$. Formally, a strategy of player i is a map

$$\sigma^i : \bigcup_{z \in \mathbf{D}} H_z^i \rightarrow (\mathbb{R}_+^L)^2 \times (\mathbb{R}_+^J)^2$$

such that $\sigma^i(h) \in A_z^i$ for all $z \in \mathbf{D}$ and for all $h \in H_z^i$.

Remark 3.1 As is well-known, strategic market games exhibit no-trade as a one-shot Nash equilibrium.²¹ As we want to prove the analog of a Folk theorem, we shall therefore need some threats that enforce the equilibrium path. Allowing for punishment phases that consist in playing the autarkic Nash one-shot equilibrium *ad libitum* would make the task rather easy. In order to prove that our result does *not* depend upon this kind of trick (hence is robust to whatever refinement would allow to get rid of the autarkic one-shot equilibrium²²), we shall focus on out-of-equilibrium strategies where players effectively trade. A second reason for not relying on the heavy hammer of autarkic Nash equilibria is that, as already said, in adverse selection problems, the market collapse has been sometimes predicted as being the unique rational consequence of differential information. We refrain from using autarky as a punishment in order to emphasize that our proof does not depend upon such a global market collapse, even as an out-of-equilibrium threat, and even though default is explicitly allowed along the equilibrium path.

3.2 Private *Interim* Beliefs

At each node z , payoffs are determined as follows: Action profile $a_z \in \mathcal{A}_z$ is played; it induces, say, x_z^i as a final allocation for player i — which is observed by i only. Notice that when entering at node z , player i may not know for sure that the current node is z . Thus, when she takes her action, she considers the expectation of her next payoffs according to her current private belief.

At each time period t , every player i updates her private belief in a Bayesian way, according to her private history. We allow for arbitrary correlation of payoffs in each state across players' utilities, endowments and assets' returns. So player i 's belief about player j 's private payoff and other higher-order beliefs are unrestricted. Let $\mathbb{P}_z^i(h_z^i) \in \Delta(\Omega)$ denote player i 's *private belief* at node z .²³ Together with actions that were played up to node z , denoted by $\sigma^i|_z$, such a probability $\mathbb{P}_z^i(h_z^i)$ induces a distribution $\mathbf{P}_z^i(h_z^i, \sigma^i|_z)$ (or $\mathbf{P}_z^i(\sigma^i|_z)$ in short) over the random characteristics of the economy to be selected after z , i.e., over $\bigcup_{z' > z} \mathcal{E}_{z'}$. In particular, it provides a distribution over i 's future payoffs which, by a slight abuse of notations, is also denoted by $\mathbf{P}_z^i(\sigma^i|_z)$.

²¹See Weyers (2004) for the elimination of this autarkic equilibrium after two rounds of elimination of dominated strategies within the extension of the Shapley-Shubik game introduced by Mertens (2003).

²²Such a refinement has been proposed, e.g. by Weyers (2004). As a consequence, Giraud and Weyers (2004) Folk theorem with complete information was already formulated so as not to rely on the autarkic threat.

²³Hereby, $\Delta(\Omega)$ is the set of all probability distributions over the finite set of states of the world.

At each node, individuals maximize her expected, discounted utility using their private *interim* belief and a common discount factor $\lambda \in (0, 1)$.²⁴ The objective function of player i is therefore of the form

$$\begin{aligned} U_{\mathbf{D}(z)}^i(x^i, \sigma^i|_z, \omega) &:= (1 - \lambda) E_{\mathbf{P}_z^i(\sigma^i|_z)} \sum_{z'=(s_{t-1}, s) \geq z} \lambda^{t-1} u_{z'}^i(x_{z'}^i) \\ &= (1 - \lambda) \sum_{z'=(s_{t-1}, s) \geq z} \lambda^{t-1} E_{\mathbf{P}_z^i(\sigma^i|_z)} \left[u_{z'}^i(x_{z'}^i) \right] \end{aligned}$$

for each node z (Given the boundedness of the utility function, the last equality is a consequence of Fubini's theorem.).

3.3 Feasible Allocations and *Interim* Individual Rationality

An *allocation* $(\bar{x}^i)_{i \in \mathcal{N}}$ is an assignment of bundles of consumption goods to every individual in the economy.

Definition 3.1 (Feasible allocation) An allocation $(\bar{x}^i)_{i \in \mathcal{N}}$ in consumption goods is said to be *feasible* if there exists a portfolio $(\bar{\varphi}^i, \bar{\theta}^i)_{i \in \mathcal{N}}$ and a price system $(\bar{p}, \bar{\pi})$ such that the following (node-wise defined) conditions are satisfied:

- *Individual budget restriction for every player i and every node $z \in \mathbf{D}$:*²⁵

$$\begin{aligned} &\sum_{\ell=1}^L \bar{p}_{z,\ell} \left(\bar{x}_{z,\ell}^i + \sum_{j=1}^J \bar{\varphi}_{z,j}^i C_{j\ell} \right) + \sum_{j=1}^J \bar{\pi}_{z,j} \left(\bar{\theta}_{z,j}^i - \bar{\varphi}_{z,j}^i \right) \\ &= \sum_{\ell=1}^L \bar{p}_{z,\ell} \left(w_{z,\ell}^i + \sum_{j=1}^J \bar{\varphi}_{z^-,j}^i C_{j\ell} \right) + \sum_{j=1}^J \left(\bar{\theta}_{z^-,j}^i - \bar{\varphi}_{z^-,j}^i \right) D_j(z) \end{aligned}$$

- *market clearing on spot markets for every good $\ell \in \mathcal{L}$ and every node:*

$$\sum_{i=1}^N \left(\bar{x}_{z,\ell}^i + \sum_{j=1}^J \bar{\varphi}_{z,j}^i C_{j\ell} \right) = \sum_{i=1}^N \left(w_{z,\ell}^i + \sum_{j=1}^J \bar{\varphi}_{z^-,j}^i C_{j\ell} \right)$$

- *market clearing on financial markets for every asset $j \in \mathcal{J}$ and every node:*

$$\sum_{i=1}^N \bar{\theta}_{z,j}^i = \sum_{i=1}^N \bar{\varphi}_{z,j}^i$$

- *and feasible trade in financial assets for every good $\ell \in \mathcal{L}$, every node and every player*

$$i: \sum_{j=1}^J \bar{\varphi}_{z,j}^i C_{j\ell} \leq w_{z,\ell}^i$$

Clearly, for every individual i , the sequence of payoffs resulting from the consumption of initial endowments is bounded from below by a constant, say \underline{u}^i . Define $\underline{u} := \min_{i \in \mathcal{N}} \underline{u}^i$. Since initial endowments are uniformly bounded, the stage-game payoff, $u_z^i(\cdot)$, induced by a feasible allocation is also uniformly bounded above by some \bar{u}^i across all action profiles, all states and all periods. Define $\bar{u} := \max_{i \in \mathcal{N}} \bar{u}^i$.

²⁴Allowing for idiosyncratic discount factors would only require notational changes.

²⁵We define $\bar{\varphi}_{z_0^-,j}^i = \bar{\theta}_{z_0^-,j}^i = 0$.

Definition 3.2 (Strongly individually rational allocation) A feasible allocation $(\bar{x}^i)_{i \in \mathcal{N}}$ is said to be *strongly individually rational* (SIR), if for all $z \in \mathbf{D}$

$$u_z^i(\bar{x}_z^i) > u_z^i(w_z^i).$$

Remark 3.2 Given a feasible allocation, notice that the SIR property holds node-wise for every state of the world $\omega \in \Omega$. Admittedly, this is a quite strong assumption. But as we do not restrict the beliefs of the players, this assumption will be useful to ensure that the feasible allocations we consider are strictly preferred to the initial endowments, independently of the players' beliefs. Note that the Folk theorem we show is partial in the sense that we consider equilibria with learning where the player's beliefs converge to the correct beliefs as given by the state of the world. There might exist other equilibria without learning.

The following Lemma says that our last two definitions generically describe a non-vacuous subset of allocations in the economy \mathcal{E}^∞ , on which, from now on, we shall focus.

Lemma 3.1 *If the initial endowments $(w_z^i)_i \gg 0$ are Pareto-inefficient in the L -good spot economy at each node $z \in \mathbf{D}$, then the economy \mathcal{E}^∞ admits a strongly individually rational and feasible (SIRF, for short) allocation.*

The next Lemma will prove useful for our main result. It shows that every SIRF allocation can be enforced by means of some adequate ("full", to be defined below) strategy. Such a strategy, however, need not fulfill any equilibrium requirement.

Definition 3.3 (Full strategy profile) A strategy profile $\sigma := (\sigma^i)_i$ is called *full* if the following holds

$$\sum_{i=1}^N q_{z,\ell}^i > 0, \quad \sum_{i=1}^N b_{z,\ell}^i > 0, \quad \sum_{i=1}^N \gamma_{z,j}^i > 0, \quad \sum_{i=1}^N \beta_{z,j}^i > 0$$

for all $\ell \in \mathcal{L}$, $j \in \mathcal{J}$, $z \in \mathbf{D}$.

Lemma 3.2 *Let $(\bar{x}^i)_{i \in \mathcal{N}}$ be a SIRF allocation. Let $(\bar{\varphi}^i, \bar{\theta}^i)_{i \in \mathcal{N}}$ and $(\bar{p}, \bar{\pi})$ be the corresponding portfolio and price system. Then $(\bar{x}^i, \bar{\varphi}^i, \bar{\theta}^i)_{i \in \mathcal{N}}$ can be implemented through the following strategy profile in the sense that, node-wise, the allocation resulting from this strategy profile is arbitrarily close to the allocation $(\bar{x}^i)_{i \in \mathcal{N}}$. Whatever being the past history, play the strategy profile consisting of the following node-wise defined actions:*

for all $z \in \mathbf{D}$, $i \in \mathcal{N}$, $\ell \in \mathcal{L}$ and $j \in \mathcal{J}$

$$\begin{aligned} q_{z,\ell}^i &= w_{z,\ell}^i + \sum_{j=1}^J \bar{\varphi}_{z^-,j}^i C_{j\ell} \\ b_{z,\ell}^i &= \bar{p}_{z,\ell} \left(\bar{x}_{z,\ell}^i + \sum_{j=1}^J \bar{\varphi}_{z,j}^i C_{j\ell} \right) \\ \gamma_{z,j}^i &= \begin{cases} \bar{\varphi}_{z,j}^i & \text{if } \sum_{i=1}^N \bar{\varphi}_{z,j}^i > 0 \\ \frac{\delta}{N} & \text{otherwise} \end{cases} \\ \beta_{z,j}^i &= \begin{cases} \bar{\pi}_{z,j} \bar{\theta}_{z,j}^i & \text{if } \sum_{i=1}^N \bar{\theta}_{z,j}^i > 0 \\ \frac{\delta}{N} & \text{otherwise} \end{cases} \end{aligned}$$

with $\delta > 0$ small. Clearly, the above strategies are full.

Remark 3.3 The reason why we do not get an exact implementation property here is that we restrict ourselves to full strategies (see Remark 3.1 supra). Indeed, were we to target an allocation that does not always require effective trade on the asset markets, then we could not target the allocation exactly. For the details we refer to the proof in the on-line Appendix A.2. This is due to the presence of the collateral constraints. Nevertheless choosing $\delta > 0$ arbitrarily small we reach an allocation that is arbitrarily close to the target allocation. A different approach is to not use full strategies on the asset markets and target the allocations exactly (as done in Brangewitz 2012, Chapter 4).

4 Learning by Trading

As already mentioned, players start with the same prior \mathbb{P} over Ω , but along the play they may (and in general they will) have different *interim* beliefs, depending upon the private information they receive.²⁶

4.1 A Leading Example

The next example exhibits a situation where one informed player trades with an uninformed one, and it is not in the interest of the informed player that her opponent learns the state of the world. Moreover, the informed player can play so as to prevent the uninformed one from learning the state of the world.

Consider an economy with $L = N = 2$. The commodities are x and y . There are two states of the world $\Omega = \{\omega_1, \omega_2\}$, and, at each period, also two states of nature $\mathcal{S} = \{s_a, s_b\}$. The transition matrix associated with ω_1 (resp. ω_2) is:

$$\omega_1 = \begin{pmatrix} 1 - \varepsilon & \varepsilon \\ \varepsilon & 1 - \varepsilon \end{pmatrix} \text{ and } \omega_2 = \begin{pmatrix} \varepsilon & 1 - \varepsilon \\ 1 - \varepsilon & \varepsilon \end{pmatrix}, \text{ with } \varepsilon > 0.$$

Suppose that $w_{z,x}^2 > w_{z,y}^1$ for every $z = (s_{t-1}, s)$ and that the individual characteristics of each trader are independent from the state of nature s , except that the upper sets of the

²⁶This is in accordance with the arguments provided by Heifetz (2006) showing that it hardly makes sense within a game-theoretic setting to assume that players start with distinct priors. Of course, Aumann's (??) theorem implies that along a play of the game it will not be common knowledge that traders have distinct *interim* beliefs.

stage-utility function of player 2 are modified by the selection of s as shown in the next Edgeworth box. Being informed of her current utility function, player 2 knows the state of nature.




Fig 2. Player 2's random preferences

The hyperbola $I_{s_b}^2$ designates the indifference curve of player 2 that crosses 2's initial endowment, $w_{s_b}^2 = w^2$ in state s_b . $I_{s_a}^2$ is the indifference curve of player 2 in state s_a crossing $w_{s_a}^2 = w^2$.

Suppose, moreover, that a central planner targets a random allocation, depending upon s that ends up at x_a whenever $s = s_a$, and x_b otherwise. The final allocation x_a is strictly preferred to $w_{s_a}^2$ by player 2 in state s_a but not in state s_b ; x_b is strictly preferred by 2 to the initial endowment in both states. Finally, suppose that player 1's prior on Ω puts a weight $1 - \varepsilon$ on ω_1 .

A single security can be traded which promises $A = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, where the first (resp. second) column stands for the return in state s_a (resp. s_b). The collateral required to sell one unit of this security is $C = (0, 1)$, i.e., one unit of commodity x .

By inspecting her own stage utility, player 2 knows the state s . Hence, after a sufficiently long time, she can know the state of the world ω with an arbitrary accuracy by using some statistical tests. From the viewpoint of player 1 on the contrary, the state of the world ω is a sunspot as it does not affect 1's characteristics. The unique way for player 1 to learn the state of nature (hence, the state of the world) requires observing the security's return in state s_a . Indeed, in state s_b , whatever being the players' strategies, the asset will default since the value of its pledged collateral always equals the value of their promise. On the contrary, in state s_a the asset will not default provided the price of commodity x is high enough — thus revealing to traders who own the asset in their portfolio (be it as sellers or as buyers) that they are in state s_a .

Suppose now that the state of nature is $\omega = \omega_2$. The prior belief for both players is such that more weight is put on ω_1 . As explained, it is in the interest of player 2 that player 1 continues to think that the state of the world is ω_1 . Assume that in $t = 1$ the state of nature s_b is chosen and in $t = 2$ it is s_a . This sequence is very likely according to ω_2 and pretty unlikely according to ω_1 . Player 2 is able to observe that s_a was chosen. Differently, player 1 cannot observe it by inspecting her private characteristics. If she knew that s_a was chosen, she would change her belief such that less weight will be put on ω_1 and therefore think that states of nature alternate between s_a and s_b . This is not in the interest of player 2 as they are supposed to coordinate one period on x_b and the other one on x_a . Therefore,

it is in player 2's interest to keep 1's opinion, i.e., to prevent 1 from updating her belief.

For this purpose, it suffices that player 2 force the default of the security in state s_a so that player 1 becomes blind after $t = 2$, i.e., $p_{z,y} > p_{z,x}$. To this end, fix the actions $(b_{z,\ell}^1, q_{z,\ell}^1)_{\ell=y,x}$ of player 1 at node $z = (s, s_b)$ on the commodity market. The following action of player 2 induces prices $p_{z,y} > p_{z,x}$ so that the security will default:

$$b_{z,x}^2 = q_{z,y}^2 = 0, q_{z,x}^2 > q_{z,y}^1 \text{ and } b_{z,y}^2 = \frac{b_{z,x}^1}{q_{z,x}^1 + q_{z,x}^2} q_{z,x}^2.$$

This action is feasible for player 2 because $w_{z,x}^2 > w_{z,y}^1$ and because her budget constraint is binding. The price of commodity y will be given by:

$$p_{z,y} = \frac{b_{z,y}^1 + \frac{b_{z,x}^1}{q_{z,x}^1 + q_{z,x}^2} q_{z,x}^2}{q_{z,y}^1} = \frac{b_{z,y}^1}{q_{z,y}^1} + \frac{q_{z,x}^2}{q_{z,y}^1} p_{z,x} > p_{z,x}.$$

The strategic challenge raised by this example lies at the heart of the proof of our main result.

4.2 The Main Result

A first lesson to be drawn from the previous example is that no learning result can be hoped for unless some restrictions are imposed on the economy.

Assumption (G). The set of L consumption goods is partitioned into two subsets, $\mathcal{L} = \mathcal{L}_a \cup \mathcal{L}_c$ with $\mathcal{L}_a \cap \mathcal{L}_c = \emptyset$. Only commodities in \mathcal{L}_c can be used as collateral, and assets' promises deliver only in commodities that belong to \mathcal{L}_a .

In other words, a commodity cannot serve both as a collateral and as a promise. This assumption is new compared to Brangewitz (2012, Chapter 4) and Giraud and Weyers (2004). That (G) cannot be dispensed with, in general, is shown by the example of the previous section. There, indeed, (G) is violated and the informed player has a mean to prevent her opponent from ever learning the state ω . By contrast, (G) enables to disentangle the prices of the bundle of goods that are promised through financial assets, and the prices of the collaterals. It will turn out that this partition of the set of commodities is sufficient to ensure that during the play of the game, a single player cannot prevent the other investors from learning the true state of the world, ω .

Given the state of the world, ω , a strategy profile, $\sigma = (\sigma^i)$, induces a unique probability distribution on the space of sequences $(w_z^i, u_z^i(\cdot), p_z(\sigma), \pi_z(\sigma), x_z^i(\sigma), A_j(z))_{i,j,z}$. Let us call this distribution, $\mathbf{P}_{\omega,\sigma}$. This is the distribution over signals from which players try to infer ω . For any two states $\omega \neq \omega'$, there must be at least some player i and some strategy profile $\sigma = (\sigma^i)_i$ such that the distributions induced by (ω, σ) and (ω', σ) over $(w_z^i, u_z^i(\cdot), p_z(\sigma), \pi_z(\sigma), x_z^i(\sigma), A_j(z))_{i,j,z}$ differ on a set of positive measure. Two states of the world that yield almost surely the same payoff, final allocation, endowment and return distributions to *every* agent and *whatever* being the strategy played, can be treated as a single state. Therefore, there is no loss of generality in assuming that a complete sequence of endowments $(w_z^i)_{i,z}$, utility functions $(u_z^i(\cdot))_{i,z}$, prices $p_z(\sigma)$ and $\pi_z(\sigma)$, final allocations $(x_z^i(\sigma))_{i,z}$, and asset returns $(A_j(z))_{j,z}$, jointly identify the state statistically for at least one well-chosen strategy profile, σ . This does not mean, however that by observing her *own private* sequence of realized individual payoffs, endowments and asset returns, a single trader is able to learn the state of the world whatever the strategy being played. Neither need prices suffice to identify the state by themselves.²⁷

²⁷When prices are interpreted as public signals, this generality contrasts with Wiseman (2011, p.8, Assumption 1) where the sole observation of public signals suffices to identify the state with no ambiguity.

The following assumption is therefore admittedly restrictive: It says that for every “reasonable” strategy profile prices, final allocations and asset returns plus individual endowments and utility functions contain *all* the relevant information about ω . Illustrations of textbook models that satisfy this assumption were given in the Introduction of the paper.

Informativeness Assumption (IA)

- (1) For any pair of nodes $(s_{t-1}, s) = z \neq z' = (s_{t-1}, s')$, any player i , and any strategy profile σ that induces an SIRF allocation at both states (s and s'), the vectors of signals, $(w_z^i, u_z^i(\cdot), p_z(\sigma), \pi_z(\sigma), x_z^i(\sigma), A_j(z))_j$ and $(w_{z'}^i, u_{z'}^i(\cdot), p_{z'}(\sigma), \pi_{z'}(\sigma), x_{z'}^i(\sigma), A_j(z'))_j$ differ.
- (2) Every ω is irreducible, aperiodic and admits an invariant measure μ_ω . Moreover, for any pair ω, ω' , if μ_ω and $\mu_{\omega'}$ are two corresponding invariant measures, then $\omega \neq \omega' \Rightarrow \mu_\omega \neq \mu_{\omega'}$.

(IA-1) says that for any strategy profile that induces an SIRF allocation, at the end of each period t , each player knows for sure at which node, $z = (s_{t-1}, s)$, she was playing, provided she can observe. Of course, this is far from sufficient in order for player i to learn ω . (IA-2) is one way of saying that two states of the world induce different distributions over states of nature in the long run. Since we are going to consider patient players, two Markov chains ω, ω' that would induce the same asymptotic distribution over signals in the long run should be identified. The last section of the paper provides some hints about how this assumption can be weakened.

In order to see that (IA) cannot be dispensed with neither, for the kind of learning result we are looking for, just consider the situation where the state of the world does not affect the players’ private characteristics, and where no asset can be traded. There is no way for players to learn the true state, so no random allocation that depends non-trivially upon the sunspot ω can be implemented.

We now turn to the (standard) perfect Bayesian equilibrium concept. Notice that in our context, a public perfect equilibrium (as defined by Fudenberg et al. (1994)) may fail to exist: First a trader’s private information influences her expectations about the future behavior of her opponents, hence about future public prices; second, the public history of prices may call for her to play a sequence of actions that she knows from her private signals will yield a payoff below that of her initial endowments (in the nodes that are going to be reached, according to her subjective beliefs). For both reasons, it may be costly for her to ignore her private information, so any best response must, in general, be a private strategy.

Definition 4.1 (Perfect Bayesian equilibrium) A pair $\left((\sigma^i)_{i \in \mathcal{N}}, (\mathbb{P}_z^i(h_z^i))_{i \in \mathcal{N}} \right)$ consisting of a feasible allocation and a system of private beliefs is a perfect Bayesian equilibrium (PBE) if

- $(\sigma^i)_{i \in \mathcal{N}}$ is strongly rational given the private beliefs $(\mathbb{P}_z^i(h_z^i))_{i \in \mathcal{N}}$, i.e., starting at any arbitrary node, given the continuation strategies of the other investors, no individual can improve her utility by unilaterally changing her strategy profile given her private beliefs $\mathbb{P}_z^i(h_z^i)$,
- and the private beliefs $(\mathbb{P}_z^i(h_z^i))_{i \in \mathcal{N}}$ are updated via Bayes rule whenever it is possible.²⁸

²⁸Due to our assumptions on the Markov chain, every state of nature is reached with a strictly positive

Our main result is that for any strategy profile that yields an allocation of commodities, assets and collaterals that is SIRF, there is a PBE in which, with arbitrarily high probability, every player achieves arbitrarily close to the allocation specified for the realized path, as long as they are patient enough. Moreover, along such an equilibrium path, every player asymptotically learns the realized state with arbitrary precision.

Theorem 1 *Under (G) and (IA), let $\varepsilon > 0$ and $(x^{*i}[\omega])_{i \in \mathcal{N}, \omega \in \Omega}$ be a SIRF allocation in consumption goods, and let \mathbb{P} be a prior belief that assigns strictly positive probability to each state of the world. Then there exists a minimal discount factor $\lambda(\mathbb{P}) < 1$ such that for all $\lambda > \lambda(\mathbb{P})$ there is a PBE that with probability at least $1 - \varepsilon$, conditional on any state ω being realized, yields a payoff vector within ε of $\left(U_{\mathbf{D}(z_0)}^i(x^{*i}, \sigma^i |_{z_0}, \omega) \right)_i$. In equilibrium, conditional on ω , each player i 's interim private belief converges to the truth: $\lim_{t \rightarrow \infty} \mathbb{P}_{z=(s_{t-1}, s)}^i(h_z^i)[\omega] = 1$ with probability 1.*

Remark 4.1 In the statement of the above theorem, note that the approximation is in terms of equilibrium payoffs and not of allocations. Here, players with incorrect beliefs learn the state of the world through coordinated experimentation. Therefore, the actual allocations during the learning phase might be far from the target allocations. Nevertheless, as the learning phase will be chosen relatively short compared to the phase where the desired allocations are targeted (using Lemma 3.2), “most of the time” not only the payoffs but the allocations as well will be close.

4.3 Proof of Theorem 2

The following sketch of the proof can serve as a lighthouse before plunging into the details.

Outline of the proof

As in Wiseman (2011), the equilibrium path uses “blocks” of $M + T$ periods each. An *equilibrium block* has a “target allocation” in commodities, denoted by

$$\left((x_z^{*i}[\omega])_i \right)_{\tau(z)=1, \dots, M+T}$$

for each state ω .²⁹ Note that by definition, there exists a corresponding portfolio

$$\left((\varphi_z^i[\omega], \theta_z^i[\omega])_i \right)_{\tau(z)=1, \dots, M+T}.$$

Within each equilibrium block, traders follow strategies that rely only on the history since the start of the block. In particular, they do not care about the history that happened before the beginning of the block. Rather, they rely on a *truncated belief*, $\mathbb{P}_{z \setminus \bar{z}}^i(h_z^i) \in \Delta(\Omega)$, defined as follows: Suppose that the block under scrutiny started at node $\bar{z} \leq z$. Given some history, h_z^i , consider the truncated history, $h_{z \setminus \bar{z}}^i$, containing only the information delivered from \bar{z} to z^- . The truncated belief, $\mathbb{P}_{z \setminus \bar{z}}^i(h_z^i) \in \Delta(\Omega)$, is the resulting updated

probability. Therefore, given the current state, Bayesian updating is always unambiguous. González-Díaz and Meléndez-Jiménez (2011) discuss the meaning of “whenever it is possible” for general extensive form games with incomplete information. In our special case their notion of a simple perfect Bayesian equilibrium coincides with usual perfect Bayesian equilibrium.

²⁹The notation of a target allocation does not explicitly contain the node (or time) at which a block starts. As all blocks are assumed to consist of the same $M + T$ periods, it can be easily found out, using the information about the node, i.e., the length of the sequence of already observed realizations of the state of nature, to which block a target allocation actually belongs.

belief starting with prior \mathbb{P} at node \bar{z} . Therefore, our model can be considered to be Markov in blocks of $M + T$ periods.

The first M periods are used in experimentation to learn the state of the world through assets' returns, initial endowments, final allocations, prices and individual stage payoffs. The most likely state, $\hat{\omega}$, is identified according to $\mathbb{P}_{z \setminus \bar{z}}^i(h_z^i)$ and, in the remaining T periods, traders choose a full action profile that yields a stream of final allocations close to the target

$$\left((x_z^{*i}[\hat{\omega}])_i \right)_{\tau(z)=M+1, \dots, M+T}$$

with utility payoffs close to

$$\left(u_z^i \left(x_z^{*i}[\hat{\omega}] \right)_i \right)_{\tau(z)=M+1, \dots, M+T}.$$

If M is large enough to identify the true state with high probability, if $T/(M + T)$ is close to one so that nearly all of the periods within the block are spent playing (close to) the target action profile, and if players are patient enough, then the expected allocation from the block when the realized state of the world is ω will be very close to the target allocation.

There are 3 types of blocks: an *equilibrium* block, a *punishment* block and a *post-deviation* block.³⁰ The initial block is equilibrium, as are all the subsequent blocks until the first deviation. If some deviation occurs during a block, it must impact prices to be profitable. Indeed, deviations that leave prices unchanged cannot influence the final allocation of goods at the end of the period, and hence cannot be profitable — a property which is specific to Shapley-Shubik games. Since prices are public signals, however, profitable deviations are immediately noticed by all the investors. Of course, a player may also want to deviate not in order to improve her current payoff but with the purpose of modifying the beliefs of her opponents. It turns out that the unique way to achieve this second goal consists in preventing the players from observing the assets' returns by provoking some default. Recall that default is not strategic in this paper. It happens as soon as the value of the collateral falls down below that of the promise (cf. equation (D)). Hence, prices must be (strategically) perturbed by the deviator in order to induce a default that was not agreed upon. We shall see in the proof how to circumvent this difficulty.

As it can be only noticed via prices, in any case a deviation remains anonymous, even when observed. Hence, punishment blocks cannot be player-specific. The next block starting after a deviation is therefore a collective punishment block. All subsequent blocks are post-deviation blocks, until a new deviation occurs. A deviation is immediately punished by switching to a punishment block.

The target allocation for each player in a punishment block, at node $z = (s_{t-1}, s)$, is made arbitrarily close to the initial endowment, w_z^i , in commodities and no-trade in financial assets. The stage payoffs of the target allocations in the post-deviation blocks are chosen to be strictly decreasing in the number of deviations so that $u_z^i \left(x_z^{i, n \text{dev}}[\omega] \right) < u_z^i \left(x_z^{i, (n-1) \text{dev}}[\omega] \right)$ for each node $z = (s_{t-1}, s)$ of the post-deviation block and each state of the world ω — where n is the number of deviations already observed.³¹ That is, the payoff to a deviator is lower than she would get in equilibrium, regardless of the state. A patient player will therefore not deviate, neither on nor off the equilibrium path, regardless of her private beliefs.

³⁰This is an important departure from the construction in Wiseman (2011), where $2N + 1$ types of blocks are at play. Both punishment and post-deviation blocks are player-specific in his setting whereas in our setting they are collective (but defined depending on the number of deviations that occurred).

³¹Of course, $u_z^i \left(x_z^{i, 0 \text{dev}}[\omega] \right) = u_z^i \left(x_z^{*i}[\omega] \right)$.

In order to understand the need for such a block decomposition, let us draw on an example (inspired from Wiseman, 2011). Suppose that the signals (endowments plus returns, prices, allocations and stage payoffs) observed by the traders strongly suggest that the state of the world is s_a ; but player's 1 private information at node z indicates state s_b more strongly. Player 1 believes that eventually everybody's belief will converge to a Dirac mass on state s_b if players continue to experiment and to learn but:

- 1) in the future, variables selected by equilibrium strategies turn out to yield the same signals in every state, so no further learning occurs. This happens, for example, if from z on, individual endowments and utility functions no more depend upon the state, $z' > z$, selected by nature, while the equilibrium strategy asks traders to trade only, say, a riskless asset whose return does not provide any information at all.
- 2) The current market belief may put so little weight on state s_b that the expected time before convergence is very long, even whenever the equilibrium path does call for further experimentation.

Further, in state s_b , the equilibrium actions specified for state s_a may yield a lower stage payoff to player 1 than her initial endowments in state s_b , i.e., than the actions designed to punish player 1 for a deviation in state s_a . And so, player 1 will deviate. In response, however, the other traders may conclude from observing unexpected prices that someone must have believed in state s_b , so the market belief may adjust toward state s_b . Then, such a deviation may be profitable for player 1 *even when her private information is consistent with state s_a* , provided the punishment profile specified in state s_b gives her a higher payoff when the actual state is s_a than does the on-path profile specified in state s_a . This can occur, again, if player 1's post-deviation payoff in s_b is higher than the final allocation induced by the equilibrium strategy profile corresponding to state s_a . So why should players different from 1 believe the anonymous deviator when she implicitly claims that the state is s_b by altering prices? Mimicking the colorful argument given by Aumann (1990, p.202) in an analogous context, players different from 1 could say: "Wait; we have a few minutes; let us think this over. Suppose that the deviator — whoever it is — doesn't trust her own claim, and so believes in state s_a . Then she would still want us to play as if we were in s_b , because that way she will get a better payoff. And of course, also if she does believe in s_b , it is better for her that we play as if we were in s_b . Thus she wants us to believe in s_b no matter what. It is as if there were no signal that 1 does not believe in s_a . So we will choose now what we would have chosen without any deviation from her." The leading example in section 4.1. provides a concrete framework where such a phenomenon can occur.

The block construction aims at circumventing this kind of stumbling rock.

Proof of Theorem 2

The proof consists in total of 8 steps.

- Step 1: Given a target allocation we construct a sequence of allocations with utility payoffs below this allocation that are used to construct a post-deviation payoff.
- Step 2: We define the δ -action profiles for the learning and punishment phase.
- Step 3: In order to start a learning phase, we define pre- M -transition action profiles.
- Step 4: To end a learning phase we define post- M -transition action profiles.
- Step 5: The block construction and the according action profiles are described.

- Step 6: The use of truncated beliefs is described and the choice of the length of the learning phase M is defined.
- Step 7: The length of the targeting period is chosen. In addition it is shown that the actual payoff is close to the target payoff.
- Step 8: It is shown that a deviation from the prescribed strategies is not profitable.

Here are the details.

Step 1.

For each state $\omega \in \Omega$, let $(x_z^{*i}[\omega])_{i,z}$ be a SIRF target allocation with stage payoffs $(v_z^{*i}[\omega])_{i,z} := (u_z^i(x_z^{*i}[\omega]))_{i,z}$. Choose a sequence of payoff vectors $((v_z^{i,n\text{dev}}[\omega])_i)_{n \in \mathbb{N}}$ that result from SIRF allocations, and such that: $v_z^{i,(n+1)\text{dev}}[\omega] < v_z^{i,n\text{dev}}[\omega]$ for every integer $n \in \mathbb{N}$, and every z — with $v_z^{i,0\text{dev}}[\omega] = v_z^{*i}[\omega]$ for each i . These utility levels will be the long-run payoffs after n deviations and can be constructed as:

$$v_z^{i,n\text{dev}}[\omega] := u_z^i(x_z^{i,n\text{dev}}[\omega]) \quad \text{with } x_z^{i,n\text{dev}}[\omega] := \rho_n x_z^{*i}[\omega] + (1 - \rho_n)w_z^i, \quad \rho_n \in (0, 1).^{32}$$

Assume that, for every $n \in \mathbb{N}$ and $z = (s_{t-1}, s) \in \mathbf{D}$:

$$0 < \varepsilon_n < v_z^{i,n\text{dev}}[\omega] - v_z^{i,(n+1)\text{dev}}[\omega] \quad (1)$$

for every player i and every state $\omega \in \Omega$. Notice that ε_n does not depend upon z , while the payoff $v_z^{i,n\text{dev}}[\omega]$ does. The sequences $(\rho_n)_n$ and $(\varepsilon_n)_n$ need to be chosen so as to converge sufficiently rapidly towards 0^+ (as $n \rightarrow +\infty$) for (1) to hold.

Step 2.

Let us now define a δ -action profile as follows.

Every player plays some action on the financial markets so that everybody gets and sells a small quantity, $\delta > 0$, of every security. Consequently, all the commodities that are eligible as collaterals will have to be partially stored. Meanwhile, on the market for consumption goods that serve as a collateral, investors bid very large quantities and offer very small quantities. As a consequence, collateral commodity prices will be large. Let us choose them sufficiently large so that there will be no default along this part of the play. And still, the quantities of commodities that are going to be effectively traded can be made arbitrarily small, as well as the quantities of collaterals they have to put aside because of their trading in securities.

The δ -actions.

Formally, for node $z = (s_{t-1}, s) \in \mathbf{D}$ in period $\tau(z) = t \in \mathbb{N}$, a δ -action is defined as follows: Let $\delta > 0$ be small. Define the actions on the goods markets by

$$b_{z,\ell}^i := \begin{cases} \bar{b}_\ell > 0 \text{ large,} & \text{for } \ell \in \mathcal{L}_c \\ \frac{\delta}{N} & \text{for } \ell \in \mathcal{L}_a \end{cases}$$

$$q_{z,\ell}^i := \frac{\delta}{N} \quad \text{for } \ell \in \mathcal{L},$$

for all $i \in \mathcal{N}$ and on the asset markets by

$$\beta_{z,j}^i := \frac{\delta}{N},$$

$$\gamma_{z,j}^i := \frac{\delta}{N}$$

³²With a slight abuse of notation we omit to denote explicitly the dependence of $v_z^{i,n\text{dev}}[\omega]$ and $x_z^{i,n\text{dev}}[\omega]$ on ρ_n . The sequence $(\rho_n)_n$ is assumed to converge to 0.

for all $j \in \mathcal{J}$, $i \in \mathcal{N}$.

It can easily be seen that these actions define feasible bids and offers and that the individual budget constraint is satisfied. The collateral requirement is equal to $\frac{\delta}{N}C_{j\ell}$, and hence as $\delta > 0$ is small, condition (F1z) and (F2z) are satisfied. Condition (F3z) is trivially satisfied as well. For the budget feasibility note that if in period $t - 1$ everybody already played a δ -action, for the current period at node z the left-hand side of the individual budget constraint ($*_z^i1$) is equal to

$$\sum_{\ell=1}^L b_{z,\ell}^i + \sum_{j=1}^J \beta_{z,j}^i = \sum_{\ell \in \mathcal{L}_c} \bar{b}_\ell + \frac{L_a \delta}{N} + \frac{J \delta}{N}$$

and the right-hand side equals

$$\begin{aligned} & \sum_{\ell=1}^L p_{z,\ell} q_{z,\ell}^i + \sum_{j=1}^J \pi_{z,j} \gamma_{z,j}^i + \sum_{j=1}^J \left(\theta_{z^-,j}^i - \varphi_{z^-,j}^i \right) D_{z,j} \\ &= \sum_{\ell \in \mathcal{L}_c} \frac{N \bar{b}_\ell}{\delta} \frac{\delta}{N} + L_a 1 \frac{\delta}{N} + J 1 \frac{\delta}{N} \\ &= \sum_{\ell \in \mathcal{L}_c} \bar{b}_\ell + \frac{L_a \delta}{N} + \frac{J \delta}{N}. \end{aligned}$$

Playing the δ -actions on asset markets every individual sells and offers the same amount of each security. Hence, net trades cancel so that no dividends will actually need to be paid.

Moreover condition ($*_z^i2$) is satisfied.

Now, what happens if player i deviates from a δ -action profile? She cannot prevent her opponents from observing their own private characteristics. Can she prevent the other players from observing the assets' returns? As she cannot prevent them from trading assets, choosing actions that induce default in all states might stop the learning process of the other players. Acting so as to decrease the price of the collateral commodities while at the same time increasing the price of the non-collateral commodities is the unique way to *cause* default. How can a single player achieve this goal? In order to decrease the price of the collateral commodities at time t , she can increase her offers on the commodity market for these goods. By doing so, she is physically constrained by her (finite) initial endowment: This is constraint (F2z). In order to be able to increase the bids for the non-collateral commodities she could first use the money from the additional sales of the collateral commodities and, second, she might have some additional dividends from asset market transactions at time $t - 1$. To satisfy the individual budget constraint ($*_z^i1$) at time $t - 1$, hence to finance the asset purchases in that very period, she needs to make some additional asset sales which are again constraint by the availability of (finite) initial endowments that need to be used to put up for the collateral: this is constraint (F2z⁻) — where z^- is the predecessor of node z . Hence, player i can neither increase the bids for non-collateral commodities arbitrarily high nor offer arbitrarily large quantities of collateral commodities. The influence on the price of player i is bounded. Thus, for each node z there exists a lower bound on the bids \bar{b}_ℓ in the δ -action profile such that if every trader bids above this bound, player i cannot induce default. From now on, a δ -action will always be understood to be such that every player's bid lies above \bar{b}_ℓ .

Step 3.

The pre-M-transition actions.

If the asset holdings are strictly positive and if players want to switch to a δ -action profile at node z , there needs to be transition period to settle the asset market obligations. Otherwise, the δ -action profile might not be budget-feasible i.e., might violate condition ($*_z^i1$).

For the pre- M -transition period at node z , define the following actions:

- on the commodity markets

$$b_{z,\ell}^i := \begin{cases} \sum_{j=1}^J \theta_{z^-,j}^i C_{j\ell} & \text{for } \ell \in \mathcal{L}_c \\ \frac{N}{\delta} & \text{for } \ell \in \mathcal{L}_a \end{cases}$$

$$q_{z,\ell}^i := \begin{cases} \sum_{j=1}^J \varphi_{z^-,j}^i C_{j\ell} & \text{for } \ell \in \mathcal{L}_c \\ \frac{\delta}{N} & \text{for } \ell \in \mathcal{L}_a \end{cases}$$

for all $\ell \in \mathcal{L}$, $i \in \mathcal{N}$.

- on the asset markets

$$\beta_{z,j}^i := \frac{\delta}{N},$$

$$\gamma_{z,j}^i := \frac{\delta}{N}$$

for all $j \in \mathcal{J}$, $i \in \mathcal{N}$.

The resulting prices are as follows:

$$p_{z,\ell} = \begin{cases} 1 & \text{for } \ell \in \mathcal{L}_c \\ \frac{N^2}{\delta^2} & \text{for } \ell \in \mathcal{L}_a \end{cases}$$

$$\pi_{z,j} = 1.$$

for $\ell \in \mathcal{L}$, $j \in \mathcal{J}$. Choose δ sufficiently small so that the prices of commodities used for the promises of assets are so large that all assets default in the transition period.

It is easy to verify that the pre- M -transition actions satisfy the feasibility constraints $(F1z)$, $(F2z)$ and $(F3z)$. For the individual budget constraint $(*_z^i 1)$ at node z we obtain for the left-hand side

$$\sum_{\ell=1}^L b_{z,\ell}^i + \sum_{j=1}^J \beta_{z,j}^i = \sum_{\ell \in \mathcal{L}_c} \sum_{j=1}^J \theta_{z^-,j}^i C_{j\ell} + \frac{L_a N}{\delta} + \frac{J\delta}{N}$$

and for the right-hand side

$$\begin{aligned} & \sum_{\ell=1}^L p_{z,\ell} q_{z,\ell}^i + \sum_{j=1}^J \pi_{z,j} \gamma_{z,j}^i + \sum_{j=1}^J \left(\theta_{z^-,j}^i - \varphi_{z^-,j}^i \right) D_{z,j} \\ &= \sum_{\ell \in \mathcal{L}_c} 1 \sum_{j=1}^J \varphi_{z^-,j}^i C_{j\ell} + L_a \frac{N^2}{\delta^2} \frac{\delta}{N} + J 1 \frac{\delta}{N} + \sum_{j=1}^J \left(\theta_{z^-,j}^i - \varphi_{z^-,j}^i \right) \left(\sum_{\ell \in \mathcal{L}_c} C_{j\ell} \right) \\ &= \frac{L_a N}{\delta} + \frac{J\delta}{N} + \sum_{\ell \in \mathcal{L}_c} \sum_{j=1}^J \theta_{z^-,j}^i C_{j\ell}. \end{aligned}$$

Moreover condition $(*_z^i 2)$ is satisfied.

From now on, unless otherwise stated, every block of δ -actions will always be preceded by the play of such transition actions.

Step 4.

The post- M -transition actions.

After the experimentation block, when a state, ω , has been identified, the investors play actions (according to Lemma 3.2) so as to target a given allocation. This target allocation

might require some holdings in certain assets which are not budget feasible given the δ -action played in the last experimentation period (e.g., a player may need to have saved much more money than she did according to the δ -action in order to finance her purchases according to the target allocation). Therefore, we add two periods of post- M -transition after the experimentation block where players can settle the asset holdings from the M -block (first post- M -transition period) and build up the necessary asset holdings for the target allocation (second post- M transition period).

Let the identified state be ω with target allocation $x_z^{*i}[\omega]$, together with actions, $\varphi_z^{*i}[\omega]$ and $\theta_z^{*i}[\omega]$, on the asset markets. The first post- M transition period is identical to a pre- M transition period (cf. *supra*). The second post- M -transition period at node z can be intuitively described as follows: People who have money from asset sales bid it on the goods markets, people who need money offer a tiny little bit of their endowment in order to get money. Commodity prices resulting from this action profile will be high, as only a little bit of commodity is offered. They turn out to be sufficiently high for every player to fulfill her budget constraint. The only point might be that some player is forced to sell a tiny little bit of her initial endowment while the collateral requirement associated with her asset sales requires her whole endowment vector to be collateralized. This would contradict (F2z). Thus, players are actually asked to sell a little bit less of assets than would be needed, were they to mimic exactly the target trades in assets. As a consequence, each player will save a small quantity of collateral that can be sold on the commodity market in order to fulfill her budget constraint. It turns out that the quantity of money lost by selling less assets can be compensated by the additional sale of commodities. More precisely,

- on the commodity markets, play:

$$b_{z,\ell}^i := \begin{cases} \frac{1}{L} \sum_{j=1}^J \pi_{z,j}^*[\omega] (\varphi_{z,j}^{*i}[\omega] - \theta_{z,j}^{*i}[\omega]) & \text{if } \sum_{j=1}^J \pi_{z,j}^*[\omega] (\varphi_{z,j}^{*i}[\omega] - \theta_{z,j}^{*i}[\omega]) \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

$$q_{z,\ell}^i := \begin{cases} -\delta \frac{1}{L} \sum_{j=1}^J \pi_{z,j}^*[\omega] (\varphi_{z,j}^{*i}[\omega] - \theta_{z,j}^{*i}[\omega]) & \text{if } \sum_{j=1}^J \pi_{z,j}^*[\omega] (\varphi_{z,j}^{*i}[\omega] - \theta_{z,j}^{*i}[\omega]) \leq 0 \\ 0 & \text{otherwise} \end{cases}$$

- on the asset markets:

$$\beta_{z,j}^i := \pi_{z,j}^*[\omega] (\theta_{z,j}^{*i}[\omega] - \eta_{z,j}^i),$$

$$\gamma_{z,j}^i := \varphi_{z,j}^{*i}[\omega] - \eta_{z,j}^i$$

for all j, i , and where $\eta_{z,j}^i := \sum_{\ell} \frac{q_{z,\ell}^i}{C_{j,\ell}}$ (with the usual convention $1/0 := 0$).

Since $x_z^{*i}[\omega]$ is feasible, the asset markets clear, $\sum_{i=1}^N \varphi_z^{*i}[\omega] = \sum_{i=1}^N \theta_z^{*i}[\omega]$. Therefore,

$$0 = \sum_{i=1}^N \sum_{j=1}^J \pi_{z,j}^*[\omega] (\varphi_{z,j}^{*i}[\omega] - \theta_{z,j}^{*i}[\omega])$$

$$= \sum_{\substack{i=1, \\ \varphi_{z,j}^{*i}[\omega] - \theta_{z,j}^{*i}[\omega] > 0}}^N \sum_{j=1}^J \pi_{z,j}^*[\omega] (\varphi_{z,j}^{*i}[\omega] - \theta_{z,j}^{*i}[\omega]) + \sum_{\substack{i=1, \\ \varphi_{z,j}^{*i}[\omega] - \theta_{z,j}^{*i}[\omega] < 0}}^N \sum_{j=1}^J \pi_{z,j}^*[\omega] (\varphi_{z,j}^{*i}[\omega] - \theta_{z,j}^{*i}[\omega])$$

Hence,

$$\sum_{\substack{i=1, \\ \varphi_{z,j}^{*i}[\omega] - \theta_{z,j}^{*i}[\omega] > 0}}^N \sum_{j=1}^J \pi_{z,j}^*[\omega] (\varphi_{z,j}^{*i}[\omega] - \theta_{z,j}^{*i}[\omega]) = - \sum_{\substack{i=1, \\ \varphi_{z,j}^{*i}[\omega] - \theta_{z,j}^{*i}[\omega] < 0}}^N \sum_{j=1}^J \pi_{z,j}^*[\omega] (\varphi_{z,j}^{*i}[\omega] - \theta_{z,j}^{*i}[\omega]).$$

The resulting prices are as follows:

$$\begin{aligned} p_{z,\ell} &= \frac{1}{\delta} \\ \pi_{z,j} &= \pi_{z,j}^*[\omega] \end{aligned}$$

for $\ell \in \mathcal{L}$, $j \in \mathcal{J}$. Choose $\delta > 0$ sufficiently small.

It is easy to verify that the transition actions satisfy the feasibility constraints ($F1z$), ($F2z$) and ($F3z$). For the asset trades note that ($F1z$) holds, as $x_z^{*i}[\omega]$ is feasible.

- Let us check whether the individual budget constraint ($*z1$) is satisfied at node z . If $\sum_{j=1}^J \pi_{z,j}^*[\omega] (\varphi_{z,j}^{*i}[\omega] - \theta_{z,j}^{*i}[\omega]) \geq 0$, we obtain for the left-hand side:

$$\sum_{\ell=1}^L b_{z,\ell}^i + \sum_{j=1}^J \beta_{z,j}^i = \sum_{j=1}^J \pi_{z,j}^*[\omega] (\varphi_{z,j}^{*i}[\omega] - \theta_{z,j}^{*i}[\omega]) + \sum_{j=1}^J \pi_{z,j}^*[\omega] \theta_{z,j}^{*i}[\omega]$$

and for the right-hand side:

$$\begin{aligned} \sum_{\ell=1}^L p_{z,\ell} q_{z,\ell}^i + \sum_{j=1}^J \pi_{z,j} \gamma_{z,j}^i + \sum_{j=1}^J (\theta_{z^-,j}^i - \varphi_{z^-,j}^i) D_{z,j} \\ = \sum_{j=1}^J \pi_{z,j}^*[\omega] \varphi_{z,j}^{*i}[\omega]. \end{aligned}$$

- If $\sum_{j=1}^J \pi_{z,j}^*[\omega] (\varphi_{z,j}^{*i}[\omega] - \theta_{z,j}^{*i}[\omega]) \leq 0$, we obtain for the left-hand side:

$$\sum_{\ell=1}^L b_{z,\ell}^i + \sum_{j=1}^J \beta_{z,j}^i = \sum_{j=1}^J \pi_{z,j}^*[\omega] (\theta_{z,j}^{*i}[\omega] - \eta_{z,j}^i),$$

and for the right-hand side:

$$\begin{aligned} \sum_{\ell=1}^L p_{z,\ell} q_{z,\ell}^i + \sum_{j=1}^J \pi_{z,j} \gamma_{z,j}^i + \sum_{j=1}^J (\theta_{z^-,j}^i - \varphi_{z^-,j}^i) D_{z,j} \\ = \sum_{j=1}^J \pi_{z,j}^*[\omega] (\theta_{z,j}^{*i}[\omega] - \varphi_{z,j}^{*i}[\omega] - \eta_{z,j}^i) + \sum_{j=1}^J \pi_{z,j}^*[\omega] \varphi_{z,j}^{*i}[\omega]. \end{aligned}$$

Thus, each budget constraint is satisfied. Finally, It is easy to see that these actions are tailored so that each player verifies the collateral constraint ($F2z$).

Step 5.

We now describe the within-block strategies.

The **equilibrium block** has length $M + T$. Suppose the true state of the world is ω . During the first M periods, play δ -actions (with a transition period if this is not the first equilibrium block of the whole play). During the first M periods of an equilibrium block, every trader is able to observe *all* assets' returns and, by combining this information with her own private initial endowments and stage payoffs, updates her prior belief, \mathbb{P} . According to (IA), by choosing M long enough, the probability that each player puts a weight larger than $1 - \varepsilon$ on the true state of the world, ω , can be made arbitrarily close to 1, whatever being $\varepsilon > 0$. More precisely, suppose that there exists a positive integer M such that, conditional

on any of the finitely many states $\omega \in \Omega$, updating the prior \mathbb{P} with the M signals that result from the δ -action profile yields a posterior truncated probability, $\mathbb{P}_{z \setminus \bar{z}}^i(h_z^i)$ ³³, for each player i that puts weight strictly greater than $1/2$ on $\{\omega\}$ with probability at least $1 - \varepsilon$. That such an integer M exists will be proven in Step 6 below.

Let $\hat{\omega}^i$ denote the state given the highest probability by player i under her own belief, $\mathbb{P}_{z \setminus \bar{z}}^i(h_z^i)$ (ties can be broken arbitrarily). Because of the choice of M (see above), the identified state $\hat{\omega}_i$ is identical across players i with probability at least $(1 - \varepsilon)^N$. Indeed, this would be the probability according to which every player, having observed her own history, will put a weight greater than $1/2$ on the true state, ω , if each history was drawn independently. Even if initial endowments and stage payoffs were probabilistically independent, the assets' returns are certainly not independent. This correlation among histories can only increase the probability above according to which players reach a consensus on the true state.

Let us denote by $\hat{\omega}$ the state on which, with probability at least $1 - \varepsilon$, players put the highest posterior probability at the end of the M -part of the block.³⁴ For the remaining T periods of the block, players start with post- M -transition actions, and then play the profile that results in a stage payoff $u_z^i(x_z^{*i}[\hat{\omega}])$ for $\tau(z) = M + 2, \dots, M + T$ in state $\hat{\omega}$. The actions are constructed using Lemma 3.2 for every node z with $\tau(z) = M + 3, \dots, M + T$. Hence, during the first M periods of an equilibrium block, individuals are learning the true state of the world. In the last $T - 2$ periods, where T is large relative to M , the target utility allocation is reached. If player i deviates unilaterally, then the equilibrium block ends immediately, and a punishment block begins in the next period. The lengths, M and T , will be chosen more precisely in steps 6 and 7.

After a deviation, a punishment phase is played, made of a certain number, P_n , of punishment blocks, each of length $M + T$, and the end of the current block. The number P_n depends on the number, n , of deviations observed. The construction of a **punishment block** is as follows. Players play *throughout* a δ -action profile as defined earlier (preceded by a transition period). This enables to learn during the punishment phase while keeping the size of net trades arbitrarily tiny. If any player unilaterally deviates from the punishment phase, then the punishment block dedicated to the first deviation ends immediately, and a new punishment phase (ending the current block and consisting in P_{n+1} blocks) begins in the next period. After the P_n punishment blocks, if no further deviation has been detected, players switch to a post-deviation block.

Play in a **post-deviation block** is divided into two parts. First, there are M periods of learning using the δ -action profiles, followed by post- M -transition actions, and finally there are $T - 2$ periods where action profiles are played, such that the target allocation after the n th deviation is reached. The target allocation in the $T - 2$ last periods of a post-deviation block consists in playing a certain sequence of SIRF allocations. Which allocations are targeted depends on the number of deviations already observed. For example, after the first deviation in the $T - 2$ last periods, the profile yielding $v_z^{i,1\text{dev}}[\hat{\omega}]$ in state $\hat{\omega}$ is played, for z with $\tau(z) = M + 3, \dots, M + T$. The first two periods after the M block are post- M -transition actions. Compared to an equilibrium block, a post-deviation block also consists of a learning phase of M periods and a target allocation in the last $T - 2$ periods. The difference is that the second sub-block does not target the equilibrium allocation but rather SIRF allocations that are strictly worse than the target allocations of the equilibrium block or the previous post-deviation block.

Step 6.

³³The current equilibrium block is supposed to start at time $\tau(\bar{z}) = \bar{t} \in \mathbb{N}$.

³⁴Ties can be broken by some arbitrary rule.

Play begins with an equilibrium block which is followed by a pre- M -transition period (for the settlement of assets' obligations) and another equilibrium block if no unilateral deviation was observed. A post-deviation- n block with no additional deviation is followed similarly by a pre- M -transition period and another post-deviation- n block. A punishment- n block (i.e., a punishment block devoted to the n th deviation) with no unilateral deviation is followed by a post-deviation- n block.

On the equilibrium path, each player's *private* belief, $\mathbb{P}_z^i(h_z^i)$, is derived by Bayesian updating the prior, \mathbb{P} , using the information of her *private* history, h_z^i . At the same time, each player computes her truncated belief, $\mathbb{P}_{z \setminus \bar{z}}^i(h_z^i)$ as defined earlier. This belief serves for the identification of the most likely state of the world, $\hat{\omega}$, according to which the allocation $x_z^{*i}[\hat{\omega}]$ is targeted during the last $T - 2$ periods of the block. By construction, $\mathbb{P}_{z \setminus \bar{z}}^i(h_z^i)$ is reset to the prior, \mathbb{P} , at the beginning of each block.

For a given ω , let the random variable, T_s^ω , be the first return time to state $s \in \mathcal{S}$:

$$T_s^\omega := \inf\{n \in \mathbb{N} \mid X_n^\omega = s\},$$

where $(X_n^\omega)_n$ stands for the stochastic process (with values in \mathcal{S}) corresponding to ω . The number

$$f_s^{(n),\omega} := \Pr(X_s^\omega = n)$$

is the probability that the process returns to state s for the first time after n steps. Since every state s is recurrent, it is easy to prove (and well-known) that the expected number of visits to s is infinite, i.e.,

$$\sum_{n \in \mathbb{N}} p_{ss}^{(n),\omega} = \pm\infty,$$

where $p_{ks}^{(n),\omega} := \Pr(X_n^\omega = s \mid X_0^\omega = k)$ for any $(k, s) \in \mathcal{S}^2$.

Since the Markov chain, ω , is irreducible and such that all the states in \mathcal{S} are positive recurrent, it admits a unique invariant measure, $\mu_\omega \in \Delta(\mathcal{S})$. The chain ω being aperiodic, the limit of the expected number, p_n , of visits of each state $s \in \mathcal{S}$ verifies:

$$\lim_{n \rightarrow +\infty} p_{ks}^{(n),\omega} = \frac{1}{E[T_s^\omega]} = \mu_\omega(s).$$

If M , the length of the experimentation block, is large enough, the probability that for every i , $\mathbb{P}_{z \setminus \bar{z}}^i(h_z^i)$ puts the maximal probability on the true state, $\{\omega\}$, at the end of the block can be made arbitrarily large: By observing the realization of their random signals, the players can observe the realization of X_ω (IA-1), hence can compute the empirical mean corresponding to the expected return time M_s of each state s . Hence, they can approximate $p_{ss}^{(n),\omega}$ with arbitrary accuracy. According to (IA-2), two different states of the world, ω, ω' , will induce different invariant measures, $\mu_\omega, \mu_{\omega'}$. Thus, for M large enough, all the players will be able to distinguish between state ω and ω' with probability at least $1 - \varepsilon$. As a consequence, all the players will learn the true state with probability at least $1 - \varepsilon$. Let us denote by M_ε the smallest such integer (whose existence was announced in Step 5 above). The crucial observation is that M_ε is independent from the discount factor λ , since it concerns only the learning process. From now on, we suppose that $M \geq M_\varepsilon$.

Step 7.

It remains to choose T large enough so that each player's welfare loss (with respect to the benchmark $v_z^{*i}[\omega]$) can be compensated by a sufficiently long targeting period of length T , provided players are sufficiently patient.

By construction of the δ -actions and by definition of the targeting actions during the T -phase of an equilibrium block, the difference between $v_z^{*i}[\omega]$ and the actual payoff that accrues to player i at node z can be made lower than ε (for δ sufficiently small). Let us denote by $U_{\mathbf{D}(z_0)}^i(\sigma^*, \omega)$ the final overall payoff induced by the equilibrium strategy, and by $U_{\mathbf{D}(z_0)}^i(x^{*i}, \omega)$, the final payoff induced by our equilibrium target allocation.³⁵

During the learning phase (of length M) and the post- M transition of two periods of each equilibrium block, the maximal stage-utility loss is \bar{u} , while during the targeting phase (of length $T - 2$), it is ε . Suppose $T - 2 = QM$ for some integer Q . One has:

$$U_{\mathbf{D}(z_0)}^i(x^{*i}, \omega) - U_{\mathbf{D}(z_0)}^i(\sigma^*, \omega) \leq (1 - \lambda) \sum_{j=0}^{+\infty} \lambda^{jQM} \left[\frac{1 - \lambda^{M+2}}{1 - \lambda} \bar{u} + \lambda^{M+2} \frac{1 - \lambda^{QM}}{1 - \lambda} \varepsilon \right],$$

where the sum of the right-hand side is taken over the sequence (indexed by j) of equilibrium blocks (of length $M + T = (Q + 1)M$). Thus,

$$U_{\mathbf{D}(z_0)}^i(x^{*i}, \omega) - U_{\mathbf{D}(z_0)}^i(\sigma^*, \omega) \leq \frac{1 - \lambda^{M+2}}{1 - \lambda^{QM}} \bar{u} + \lambda^{M+2} \varepsilon.$$

For every $\varepsilon > 0$, there exists some $Q_{\varepsilon, \lambda}$ large enough and some λ_ε close enough to 1, so the right-hand side of the last inequality is lower than ε . From now on, we assume that $Q \geq Q_\varepsilon$ and $\lambda \geq \lambda_\varepsilon$.

Therefore, along the equilibrium path, players learn the true state with probability at least $1 - \varepsilon$ and their final payoff will be within ε of the benchmark. It follows that a patient player prefers not to deviate even if the truncated belief after a sequence of misleading signals calls for an action profile that she thinks will give her a very low payoff for the duration of the current block: At the start of the next block, the pseudo-belief will revert to the prior, and with high likelihood, experimentation in the next blocks will reveal the true state of the world and enable the other players to provide her with the equilibrium payoff or to effectively punish her in case of deviation.

Actions may reveal a piece of information about a player's private payoffs. For instance, by deviating, player i may induce a final allocation for player j different from the one that is prescribed at equilibrium. This different allocation may in turn provide j with some information in terms stage payoff that was out of scope with the equilibrium allocation. And even during a punishment phase, a deviator might be tempted to keep talking with her opponents through the manipulation of their commodity allocations. Nevertheless, (IA) implies that as long as they still observe every asset's return, all the players will learn the true state with arbitrary precision *whatever* being their stream of stage payoffs.³⁶ By manipulating allocations (hence stage payoffs), a player cannot prevent her opponents from eventually learning the true state of the world, ω .

Step 8.

It remains to choose P_n (the number of punishment blocks after n deviations) large enough so that no player has any incentive to deviate, neither on the equilibrium path nor off this path, whatever her private belief about ω or her higher order beliefs (about others' beliefs). For this purpose, we need to guarantee that a post-deviation long-run discounted payoff never exceeds the equilibrium long-run discounted payoff. Suppose that the deviation occurs at node $z' = (s'_{t-1}, s')$, that it is the $(n + 1)$ th deviation observed during the play

³⁵The slight abuse of notations in the arguments of the overall utility should not create any confusion.

³⁶In other words, players need to be able to observe the sequence of stage payoffs resulting from *some* SIRF allocation, plus asset returns and initial endowments. For a patient player, the choice of the *particular* sequence of SIRF allocations is irrelevant.

and that there are no further deviations at later nodes. It will at most yield \bar{u}^i to player i . Then, the post-deviation payoff can be made ε -close to the following maximum:

$$\begin{aligned}
 & (1 - \lambda)\lambda^{t'-1} \left[\bar{u}^i + \sum_{t=2}^{M+T} \lambda^{t-1} E_{\mathbf{P}_{z'}^i(\sigma|_{z'})} [u_z^i(w_z^i)] \right. \\
 & + \sum_{k=1}^{P_{n+1}} \left[\sum_{t=k(M+T)+1}^{(k+1)(M+T)} \lambda^{t-1} E_{\mathbf{P}_{z'}^i(\sigma|_{z'})} [u_z^i(w_z^i)] \right. \\
 & + \sum_{k \geq P_{n+1}+1} \left[\sum_{t=k(M+T)+1}^{k(M+T)+M} \lambda^{t-1} E_{\mathbf{P}_{z'}^i(\sigma|_{z'})} [u_z^i(w_z^i)] \right. \\
 & + \sum_{t=k(M+T)+M+1}^{k(M+T)+M+2} \lambda^{t-1} \bar{u} \\
 & \left. \left. + \sum_{t=k(M+T)+M+3}^{(k+1)(M+T)} \lambda^{t-1} E_{\mathbf{P}_{z'}^i(\sigma|_{z'})} [v_z^{i,(n+1)\text{dev}}[\omega]] \right] \right]. \tag{2}
 \end{aligned}$$

Indeed, the long-run discounted payoff after a deviation consists once of a (maybe) very high payoff from deviating, then the payoff from a punishment during the current block plus P_{n+1} punishment blocks lasting $M+T$ periods, and finally the payoff from succeeding post-deviation- $(n+1)$ blocks of $M+T$ periods including possibly a very high payoff in the post- M -transition period. On the other hand, since no deviator can prevent her opponents from learning the state of the world with arbitrary precision (even during the punishment phase and whatever being the behavior of the deviator), the reward payoff, $E_{\mathbf{P}_{z'}^i(\sigma|_{z'})} [v_z^{i,(n+1)\text{dev}}[\hat{\omega}]]$, computed with the most likely state, $\hat{\omega}$ (according to the players' truncated belief), can also be made arbitrarily close to $E_{\mathbf{P}_{z'}^i(\sigma|_{z'})} [v_z^{i,(n+1)\text{dev}}[\omega]]$.

By contrast, if the $(n+1)$ th deviation did not take place, i 's long-run discounted payoff would consist in the payoff from post-deviation- n blocks of $M+T$ periods. Therefore, it would be arbitrarily close to:

$$\begin{aligned}
 & (1 - \lambda)\lambda^{t'-1} \sum_{k \geq 0} \left[\sum_{t=k(M+T)+1}^{k(M+T)+M} \lambda^{t-1} E_{\mathbf{P}_{z'}^i(\sigma|_{z'})} [u_z^i(w_z^i)] \right. \\
 & \left. + \sum_{t=k(M+T)+M+3}^{(k+1)(M+T)} \lambda^{t-1} E_{\mathbf{P}_{z'}^i(\sigma|_{z'})} [v_z^{i,n\text{dev}}[\omega]] \right]. \tag{3}
 \end{aligned}$$

Note that here we assumed a payoff of 0 in the post- M -transition periods.

In order to check whether the difference (3) - (2) is positive, all we need is to ensure that:

$$\begin{aligned}
 & (1 - \lambda)\lambda^{t'-1} \left[\underline{u} - \bar{u} + \sum_{k=1}^{P_{n+1}} \sum_{t=k(M+T)+M+3}^{(k+1)(M+T)} \lambda^{t-1} E_{\mathbf{P}_{z'}^i(\sigma|_{z'})} [v_z^{i,n\text{dev}}[\omega] - u_z^i(w_z^i)] \right. \\
 & \left. + \varepsilon_n \sum_{k \geq P_{n+1}+1} \left[\sum_{t=k(M+T)+M+3}^{(k+1)(M+T)} \lambda^{t-1} \right] - \bar{u} \sum_{k \geq 0} \left[\sum_{t=k(M+T)+M+3}^{k(M+T)+M+2} \lambda^{t-1} \right] \right] > 0
 \end{aligned}$$

Note that since $v_z^{i,n\text{dev}}[\omega]$ results from a SIRF allocation, we have

$$E_{\mathbf{P}_{z'}^i(\sigma|_{z'})} [v_z^{i,n\text{dev}}[\omega] - u_z^i(w_z^i)] > 0$$

for every node $z = (s_{t-1}, s)$, and every individual i . Let us define

$$g_z := \min_{i \in \mathcal{N}} E_{\mathbf{P}_{z'}^i(\sigma|_{z'})} [v_z^{i, \text{ndev}}[\omega] - u_z^i(w_z^i)].$$

It is sufficient to require that:

$$\begin{aligned} & (1 - \lambda) \lambda^{t'-1} \left[\underline{u} - \bar{u} + \sum_{k=1}^{P_{n+1}} \left[\sum_{t=k(M+T)+M+3}^{(k+1)(M+T)} \lambda^{t-1} g_z \right] \right. \\ & + \varepsilon_n \frac{(1 - \lambda^{T-2})}{(1 - \lambda)(1 - \lambda^{M+T})} \lambda^{(P_{n+1}+1)(M+T)+M+2} \\ & \left. - \bar{u} \frac{(1 - \lambda^2)}{(1 - \lambda)(1 - \lambda^{M+T})} \lambda^M \right] > 0. \end{aligned}$$

It is easy to see that whatever being the distance, $(1 - \lambda)\underline{u} - \left((1 - \lambda) + \frac{(1 - \lambda^2)}{(1 - \lambda^{M+T})} \lambda^M \right) \bar{u}$, and that for every $\varepsilon_n > 0$ and every $\lambda \geq \lambda_\varepsilon$, there exists some integer $P_{n+1}^{\lambda, \varepsilon}$ big enough so that this last inequality is satisfied.

Suppose a deviator keeps deviating. While being punished by δ -actions, the most she can grasp is δ units of each commodity in each period. Immediately a new punishment starts with a punishment phase at least as long as the one before. As the reward in the post-deviation block declines, continuing deviating becomes even less attractive as the payoff in equation (2).

This completes the proof that it is in no player's interest to deviate from the prescribed equilibrium strategy, be it on the equilibrium path (i.e., whenever no deviation already occurred) or out of the equilibrium path (i.e., after a deviation occurred), provided: $M \geq M_\varepsilon$, $\lambda \geq \lambda_\varepsilon$, $T = QM$ with $Q \geq Q_\varepsilon$, and $\forall n, P_n \geq P_n^{\lambda, \varepsilon}$. □

5 A Nonconvergence Result

In this section we apply the preceding result to study the efficiency properties of the Nash equilibria of the infinite horizon strategic market game where each player becomes negligible. Postlewaite and Schmeidler (1978) show for the truncated game with finite horizon that if there are sufficiently many traders and if aggregate initial endowments are large enough relative to the number of traders, then any full Nash equilibrium is η -efficient, with $\eta > 0$ arbitrarily small. Contrary to what is obtained in Postlewaite and Schmeidler (1978), for the one-shot game — but in accordance with the result obtained in Giraud and Weyers (2004) for the finite horizon case-complete information case — here we do not obtain a convergence to “almost” efficient equilibria. Quite the contrary, we easily deduce from our Folk theorem that whatever being the number of replica per player, the set of SIRF allocations remains approximately reachable by (type-symmetric) Bayesian perfect equilibria. Obviously, this is due to the fact that however small the weight of a player is, as long as it is non-zero, this player can influence prices so that the monitoring of prices reveals a relevant information across time — hence opens the door to the whole machinery underlying our preceding result.

Definition 5.1 For $\eta > 0$, a SIRF allocation $(x^i)_{i \in \mathcal{N}}$ is η -efficient if there does not exist any SIRF allocation that Pareto dominates $(x^i)_{i \in \mathcal{N}}$ in any fictitious economy with the same utility functions and in which the aggregate initial endowment vectors are at most $(1 - \eta) \sum_i w^i$ — where $(w^i)_{i \in \mathcal{N}}$ are the initial endowments in the original economy.

For the sake of completeness, let us recall the following result stated in the one-shot game associated to the one-period Arrow-Debreu economy, $\mathcal{E}_z = (u^i, w^i)_i$ (at node $z \in \mathbf{D}$).

Approximate Efficiency Theorem. (Postlewaite and Schmeidler 1978) *For any positive number α, β, η any allocation resulting from a full Nash equilibrium in an economy with $u_z^i(\cdot)$ and w_z^i is η -efficient, where the initial endowments $w_z^i \leq (\beta, \dots, \beta)$ for all individuals $i \in \mathcal{N}$ and $\sum_{i=1}^N w_z^i > (N\alpha, \dots, N\alpha)$ and $N > \frac{16L\beta}{\alpha\eta^2}$*

In order to be able to compare this result with the properties of our infinite-horizon economy, recall that whenever initial endowments are uniformly bounded, the commodity space of the infinite-horizon economy, \mathcal{E}^∞ , can be viewed as the space, $\ell^\infty(\mathbf{D} \times L)$, of uniformly bounded stochastic processes defined over the event tree, \mathbf{D} , and taking values in \mathbb{R}^L . Let us endow this commodity space with the classical Mackey topology, $\tau(\ell^\infty, \ell^1)$.³⁷ By contrast with the Approximate Efficiency Theorem just recalled, we get the following corollary of Theorem 1:

Corollary 5.1 *Suppose that (G) and (IA) are in force. Take $\eta > 0$ and an integer N , arbitrary. Whatever being the number, N , of players, if they are sufficiently patient, if $(w^i)_{i \in \mathcal{N}}$ is not in the $\tau(\ell^\infty, \ell^1)$ -closure of the set of η -efficient allocations, then there exists a non-empty, $\tau(\ell^\infty, \ell^1)$ -open subset of $\ell^\infty(\mathbf{D})$ consisting of allocations that are not η -efficient and that can nevertheless be approximated by a PBE (in the sense of Theorem 1).*

We begin with:

Lemma 5.1 *For any number, N , of players,*

- (i) *if $(w^i)_{i \in \mathcal{N}}$ is η -efficient, then every SIRF allocation is η -efficient.*
- (ii) *if $(w^i)_{i \in \mathcal{N}}$ is not in the $\tau(\ell^\infty, \ell^1)$ -closure of the set of η -efficient allocations, then the set of SIRF allocations contains a $\tau(\ell^\infty, \ell^1)$ -open subset that is not in the $\tau(\ell^\infty, \ell^1)$ -closure of the set of η -efficient allocations.*

Proof of Lemma 5.1.

Let us first consider the economy, \mathcal{E}^T , obtained after having truncated \mathcal{E} at T .

(i) Suppose $(w_z^i)_{i \in \mathcal{N}, \tau(z) \leq T}$ is η -efficient in \mathcal{E}^T . Let $(x_z^i)_{i \in \mathcal{N}, \tau(z) \leq T}$ be a SIRF allocation that is not η -efficient. Then by the definition of η -efficiency there exists $y_z = (y_z^i)_{i \in \mathcal{N}, \tau(z) \leq T}$ such that

$$u_z^i(y_z^i) \geq u^i(x_z^i)$$

$$\sum_{i=1}^N y_z^i \leq (1 - \eta) \sum_{i=1}^N x_z^i = (1 - \eta) \sum_{i=1}^N w_z^i$$

for at least one $z \in \mathbf{D}$. Since x_z^i is SIR we have $u_z^i(x_z^i) > u_z^i(w_z^i)$. Hence, $u_z^i(y_z^i) > u_z^i(w_z^i)$ and $\sum_{i=1}^N y_z^i \leq (1 - \eta) \sum_{i=1}^N w_z^i$ at $z \in \mathbf{D}$. Therefore, leaving the allocations at nodes other than z unchanged shows that this is in contradiction to the assumption that w^i was η -efficient.

(ii) Suppose $(w_z^i)_{i \in \mathcal{N}, \tau(z) \leq T}$ is not η -efficient. Take $(y_z^i)_{i \in \mathcal{N}, \tau(z) \leq T}$ feasible that strictly Pareto-dominates the initial endowments $(w_z^i)_{i \in \mathcal{N}, \tau(z) \leq T}$, as in the proof of Lemma 3.1, and for $z \in \mathbf{D}$ with $\tau(z) \leq T$ let $y_{k,z}^i = \frac{1}{k} y_z^i + \frac{k-1}{k} w_z^i$ for all $i \in \mathcal{N}$. Then $(y_{k,z}^i)_{i \in \mathcal{N}, \tau(z) \leq T}$ is SIRF

³⁷That is, the strongest topology such that $\ell^1(\mathbf{D} \times L)$ is the topological dual of $\ell^\infty(\mathbf{D} \times L)$, see, e.g., Magill and Quinzii (1994).

for all k . Because $(w_z^i)_{i \in \mathcal{N}, \tau(z) \leq T}$ is not in the closure of the set of η -efficient allocations, there exists a k_0 such that for every $k > k_0$ we have that $(y_{k,z}^i)_{i \in \mathcal{N}, \tau(z) \leq T}$ is not in the closure of the set of η -efficient allocations. Because of the strict concavity of preferences we have $u_z^i(y_{k,z}^i) > u_z^i(w_z^i)$ for all $i \in \mathcal{N}$. Thus, for all nodes $z \in \mathbf{D}$ with $\tau(z) \leq T$ we have that $(y_{k,z}^i)_{i \in \mathcal{N}}$ is in the interior of the individual rational and feasible allocations of the spot economy. Since preferences are strictly monotonic, for every feasible allocation belonging to some neighborhood of $(y_{k,z}^i)_{i \in \mathcal{N}}$ there is a hyperplane through this allocation and the initial endowment $(w_z^i)_{i \in \mathcal{N}}$ with a strictly positive price vector. Thus, there exists a r_z such that all allocations in the ball, $B\left((y_{k,z}^i)_{i \in \mathcal{N}}, r_z\right)$, with radius r_z at center $(y_{k,z}^i)_{i \in \mathcal{N}}$, are feasible and individually rational. For r_z sufficiently small the set $B\left((y_{k,z}^i)_{i \in \mathcal{N}}, r_z\right)$ is not contained in the closure of the η -efficient allocations. Since this argument holds for all $z \in \mathbf{D}$ with $\tau(z) \leq T$ we have found an open set, $B = \prod_{z \in \mathbf{D}, \tau(z) \leq T} B\left((y_{k,z}^i)_{i \in \mathcal{N}}, r_z\right)$, that is SIRF and full-dimensional. For $\max_{z \in \mathbf{D}, \tau(z) \leq T} r_z$ sufficiently small, the set B is not contained in the closure of the η -efficient allocations.

Letting $T \rightarrow +\infty$, it follows from the Mackey-continuity of the players' preferences (under our general assumptions, see ?) that, if $(w^i)_{i \in \mathcal{N}}$ is not in the $\tau(\ell^\infty, \ell^1)$ -closure of the set of η -efficient (stream of) allocations, then there exists a non-empty, $\tau(\ell^\infty, \ell^1)$ -open subset of $\ell^\infty(\mathbf{D})$ that is not in the $\tau(\ell^\infty, \ell^1)$ -closure of the set of η -efficient allocations. \square

Proof of Corollary 5.1. It follows from Theorem 1 and Lemma 5.1 that whatever being the number N , if players are sufficiently patient, then under (G) and (IA), if $(w^i)_{i \in \mathcal{N}}$ is not in the $\tau(\ell^\infty, \ell^1)$ -closure of the set of η -efficient allocations, there exists a non-empty, $\tau(\ell^\infty, \ell^1)$ -open subset of $\ell^\infty(\mathbf{D})$ that is not in the $\tau(\ell^\infty, \ell^1)$ -closure of the set of η -efficient allocations and that can nevertheless be approximated by a PBE (in the sense of Theorem 1). \square

6 Concluding Comments

In this paper we investigated the general properties of perfect Bayesian equilibria in imperfectly competitive environments with incomplete information. We proved that adding collateral constraints within the rules of trading has an ambiguous effect. Collateral constraints limit the extent to which agents can pledge their future wealth and ensure that players with incorrect beliefs never lose so much as to be driven out of the market. Consequently all agents, regardless of their beliefs, survive in the long run and continue to trade, possibly on the basis of those heterogeneous beliefs. Cao (2011) showed that the presence of heterogeneous beliefs together with collateral requirements lead to additional leverage and asset price volatility (relative to a model with homogeneous beliefs or relative to the complete markets economy). By contrast, here we show that virtually any kind of volatility (in the SIRF allocation of commodities and assets) is compatible with the imperfectly competitive equilibrium condition. When investors can influence prices, indeed, traders with incorrect beliefs can strategically *learn* the state of the world. We thus provided a partial characterization of learning equilibria, at the end of which no player shares incorrect beliefs — not because they were eliminated from the market (although default is possible at equilibrium) but because they have taken time to update their prior belief. The striking point is that our (partial) Folk theorem provides us with a wide range of equilibria, many of them being first-best efficient, many others being dominated. This, at the very least, suggests that considerable care is necessary in invoking the impact of collateral regulation on the inefficiency of equilibria with asymmetric information.

In terms of future research, we see the following directions: First, the robustness of our result with respect to the uncertainty setup is unclear (at least to us). Exploring the

situation of a reducible Markov chain (or, possibly, a non-Markovian stochastic process) would certainly provide new challenges. Second, it is well-known that several infinite-horizon games with incomplete information have no Bayes-Nash equilibria (not to speak about perfect equilibria) once the discounted utility function is replaced by some limit of the arithmetic mean. Thus, there is, in general, a fundamental discontinuity of the equilibrium set as the discount factor goes to one. Does this discontinuity pervade in the specific situation of a strategic market game? Next, the financial friction adopted in this paper consisted in imposing a collateral constraint together with short sales. Many other financial frictions have been studied in the literature, such as pledgeable income models (?), models of private information and costly state verification (?) or self-enforcing debt constraints (Kehoe and Levine (1993)). Each of these alternatives can be tested to improve the robustness of our result with respect to the institutional specification of the financial frictions at play. Finally, Shapley-Shubik games are but *one* type of market games. Many others can be thought of (see Giraud (2003)), which all require a specific treatment.

Most importantly, however, our result suggests that the conclusion in Cao (2011) might be partly due to the narrow (though perfectly standard) definition of perfect competition as price-taking behavior adopted in this paper. To recover perfect competition from the imperfectly competitive setting studied here, a familiar way consists in increasing the number of replica per player. This is where we end up with a surprise: We get that at the limit, there are still plenty of perfect Bayesian equilibria exhibiting a large variety of efficiency properties (although each individual is asymptotically negligible). The limit benchmark obtained by letting the weight of each price maker shrink to zero does not lead to the same conclusion as a direct price-taking assumption. This is clearly due to the fact that whenever perfect competition is understood as an asymptotic notion, the monitoring (of private signals and public prices) sharply changes the picture of equilibria. To put it differently, the asymptotic approach of perfect competition reveals that this notion not only rests on some price-taking assumption but also on some (hidden) no-monitoring restriction. This suggests that despite the considerable literature devoted to its foundation, the very concept of perfect competition itself deserves further investigation.

References

- (1956). The analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1):1–8. 4
- (2010). Dynamic sender-receiver games. *Preprint*. 3
- Acemoglu, D. and Simsek, A. (2010). Moral hazard and efficiency in general equilibrium with anonymous trading. Working Paper 10-8, MIT Department of Economics. Available at SSRN: <http://ssrn.com/abstract=1597907>. 6
- Akerlof, G. A. (1970). The market for “lemons”: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500. 6
- Amir, R., Sahi, S., Shubik, M., and Yao, S. (1990). A strategic market game with complete markets. *Journal of Economic Theory*, 51(1):126–143. 9
- Araujo, A., Pascoa, M. R., and Torres-Martinez, J. P. (2002). Collateral avoids Ponzi schemes in incomplete markets. *Econometrica*, 70(4):1613–1638. 5
- Arnott, R. and Stiglitz, J. E. (1986). Moral hazard and optimal commodity taxation. *Journal of Public Economics*, 29(1):1–24. 6

- Arnott, R. and Stiglitz, J. E. (1990). The welfare economics of moral hazard. In *H. Louberge (editor) Risk, Information and Insurance: Essays in the Memory of Karl H. Borch*, pages 91–122. Kluwer. 6
- Arnott, R. and Stiglitz, J. E. (1991). Moral hazard and nonmarket institutions: Dysfunctional crowding out of peer monitoring? *The American Economic Review*, 81(1):179–190. 6
- Aumann, R. (1990). Nash equilibria are not self-enforcing. In *J. J. Gabszewicz, J.-F. Richard, and L. A. Wolsey, eds., Economic Decision-Making: Games, Econometrics and Optimisation*, pages 201–206. Amsterdam: Elsevier. 21
- Blume, L. and Easley, D. (2006). If you're so smart, why aren't you rich? belief selection in complete and incomplete markets. *Econometrica*, 74(4):929–966. 2
- Brangewitz, S. (2012). *Coalitional and Strategic Market Games*. PhD thesis, Bielefeld University. 1, 2, 15, 17
- Cao, D. (2011). Collateral shortages, asset price and investment volatility with heterogeneous beliefs. Unpublished. 2, 6, 33, 34
- De Meyer, B. and Saley, H. M. (2003). On the strategic origin of brownian motion in finance. *International Journal of Game Theory*, 31:285–319. 2
- Debreu, G. (1952). A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences of the United States of America*, 38(10):886–893. 11
- Dubey, P. and Geanakoplos, J. (2003). Monetary equilibrium with missing markets. *Journal of Mathematical Economics*, 39(5-6):585–618. 10
- Duffie, D. and Shafer, W. (1985). Equilibrium in incomplete markets: I : A basic model of generic existence. *Journal of Mathematical Economics*, 14(3):285–300. 4
- Escobar, J. and Toikka, J. (2012). Efficiency in games with markov private information. *Preprint*. 2, 3
- Facchinei, F. and Kanzow, C. (2010). Generalized Nash equilibrium problems. *Annals of Operations Research*, 175(1):177–211. 11
- Florenzano, M. and Gourdel, P. (1996). Incomplete markets in infinite horizon: Debt constraints versus node prices. *Mathematical Finance*, 6(2):167–196. 5
- Friedman, M. (1953). *Essays in Positive Economics*. University of Chicago Press. 3
- Fudenberg, D., Levine, D., and Maskin, E. (1994). The folk theorem with imperfect public information. *Econometrica*, 62(5):997–1039. 2, 18
- Geanakoplos, J. and Zame, W. R. (2007). Collateralized asset markets. *UCLA Working Paper*. 3
- Giraud, G. (2003). Strategic market games: an introduction. *Journal of Mathematical Economics*, 39:355–375. 3, 34
- Giraud, G. and Weyers, S. (2004). Strategic market games with a finite horizon and incomplete markets. *Economic Theory*, 23:467–491. 2, 9, 12, 17, 31, 37, 38
- Golosov, Skreta, L., Tsyvinski, V., and Wilson, A. (2009). Dynamic strategic information transmission. *Preprint*. 2

- González-Díaz, J. and Meléndez-Jiménez, M. A. (2011). On the notion of perfect Bayesian equilibrium. Unpublished. 19
- Greenwald, B. C. and Stiglitz, J. E. (1984). Externalities and economies with imperfect information and incomplete markets. *Quarterly Journal of Economics*, 101(2):229–264. 6
- Harker, P. T. (1991). Generalized Nash games and quasi-variational inequalities. *European Journal of Operational Research*, 54(1):81–94. 11
- Hart, O. D. (1975). On the optimality of equilibrium when the market structure is incomplete. *Journal of Economic Theory*, 11(3):418–443. 10
- Hart, S. (1985). Nonzero-sum two-person repeated games with incomplete information. *Mathematics of Operations Research*, 10(1):pp. 117–153. 2
- He, Z. and Xiong, W. (2010). Equilibrium debt financing. Unpublished. 2
- Heifetz, A. (2006). The positive foundation of the common prior assumption. *Games and Economic Behavior*, 56(1):105–120. 15
- Kehoe, T. J. and Levine, D. K. (1993). Debt-constrained asset markets. *The Review of Economic Studies*, 60(4):865–888. 34
- Kubler, F. and Schmedders, K. (2003). Stationary equilibria in asset-pricing models with incomplete markets and collateral. *Econometrica*, 71(6):pp. 1767–1793. 6
- Magill, M. and Quinzii, M. (1994). Infinite horizon incomplete markets. *Econometrica*, 62(4):853–880. 7, 32
- Mertens, J. F. (2003). The limit-price mechanism. *Journal of Mathematical Economics*, 39:433–528. 12
- O’Connell, S. A. and Zeldes, S. P. (1988). Rational Ponzi games. *International Economic Review*, 29(3):431–450. 5
- Peck, J., Shell, K., and Spear, S. E. (1992). The market game: Existence and structure of equilibrium. *Journal of Mathematical Economics*, 21:271–299. 9
- Postlewaite, A. and Schmeidler, D. (1978). Approximate efficiency of non-Walrasian Nash equilibria. *Econometrica*, 46(1):127–135. 9, 31, 32
- Prescott, E. C. and Townsend, R. M. (1984a). General competitive analysis in an economy with private information. *International Economic Review*, 25(1):1–20. 6
- Prescott, E. C. and Townsend, R. M. (1984b). Pareto optima and competitive equilibria with adverse selection and moral hazard. *Econometrica*, 52(1):21–46. 6
- Radner, R. (1979). Rational expectations equilibrium: Generic existence and the information revealed by prices. *Econometrica*, 47(3):655–678. 4
- Rahi, R. (1995). Optimal incomplete markets with asymmetric information. *Journal of Economic Theory*, 65(1):171–197. 5
- Rothschild, M. (1974). A two-armed bandit theory of market pricing. *Journal of Economic Theory*, 9(2):185–202. 4

- Rothschild, M. and Stiglitz, J. (1976). Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, 90(4):629–649. 6
- Shapley, L. S. and Shubik, M. (1977). Trade using one commodity as a means of payment. *Journal of Political Economy*, 85:937–968. 9
- Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–374. 6
- Thomas, C. P. (1995). The role of fiscal policy in an incomplete markets framework. *The Review of Economic Studies*, 62(3):449–468. 2
- Weyers, S. (2004). Convergence to competitive equilibria and elimination of no-trade (in a strategic market game with limit prices). *Journal of Mathematical Economics*, 40(8):903–922. 12
- Wiseman, T. (2011). A partial folk theorem for games with private learning. Forthcoming in *Theoretical Economics*. 1, 3, 11, 17, 19, 20, 21
- Zhu, H. (2011). Do dark pools harm price discovery? Available at SSRN: <http://ssrn.com/abstract=1712173>. 5

A Appendix

A.1 Proof of Lemma 3.1

To show Lemma 3.1 for our model, we modify the proof of Giraud and Weyers (2004) slightly.

Proof. Fix a node $z \in \mathbf{D}$. Since the allocation of initial endowments $(w_z^i)_i$ are Pareto-inefficient in the L -good spot economy, there exists a feasible allocation $(\bar{x}_z^i)_i$ that Pareto dominates $(w_z^i)_i$ and satisfies for every good $\ell \in \mathcal{L}$

$$\sum_{i=1}^N \bar{x}_{z,\ell}^i = \sum_{i=1}^N w_{z,\ell}^i.$$

By the strict monotonicity of the preferences, there exists a feasible allocation $(\bar{x}_z^i)_i$ such that

$$u_z^i(\bar{x}_z^i) > u_z^i(w_z^i) \quad i = 1, \dots, N$$

and

$$\sum_{i=1}^N \bar{x}_{z,\ell}^i = \sum_{i=1}^N w_{z,\ell}^i.$$

Since the utility functions are strictly increasing, there exists a hyperplane containing $(\bar{x}_z^i)_i$ and $(w_z^i)_i$ with a strictly positive price vector p_z . Thus, the individual budget restriction

$$p_z \cdot \bar{x}_z^i = p_z \cdot w_z^i$$

is satisfied and furthermore

$$u_z^i(\bar{x}_z^i) > u_z^i(\bar{x}_z^i) \geq u_z^i(w_z^i)$$

for all $i \in \mathcal{N}$. This argument holds for all $z \in \mathbf{D}$, which proves the claim. \square

A.2 Proof of Lemma 3.2

To show Lemma 3.2 for our model, we modify the proof of Giraud and Weyers (2004).

Proof. Since $(\bar{x}^i)_{i \in \mathcal{N}}$ is a feasible allocation there exist $(\bar{\varphi}^i, \bar{\theta}^i)_{i \in \mathcal{N}}$ such that the asset markets clear at every node $z \in \mathbf{D}$. For all $j \in \mathcal{J}$ we have

$$\sum_{i=1}^N \bar{\theta}_{z,j}^i = \sum_{i=1}^N \bar{\varphi}_{z,j}^i.$$

Therefore, if $\sum_{i=1}^N \bar{\theta}_{z,j}^i = 0$, then $\sum_{i=1}^N \bar{\varphi}_{z,j}^i = 0$ and vice versa.

Using the market clearing condition on the goods markets we obtain from the definition of the actions

$$\begin{aligned} \sum_{i=1}^N q_{z,\ell}^i &= \sum_{i=1}^N \left(w_{z,\ell}^i + \sum_{j=1}^J \bar{\varphi}_{z,j}^i C_{j\ell} \right) = \sum_{i=1}^N \left(\bar{x}_{z,\ell}^i + \sum_{j=1}^J \bar{\varphi}_{z,j}^i C_{j\ell} \right), \\ \sum_{i=1}^N b_{z,\ell}^i &= p_{z,\ell} \sum_{i=1}^N \left(\bar{x}_{z,\ell}^i + \sum_{j=1}^J \bar{\varphi}_{z,j}^i C_{j\ell} \right). \end{aligned}$$

Hence,

$$\bar{p}_{z,\ell} = \frac{\sum_{i=1}^n b_{z,\ell}^i}{\sum_{i=1}^n q_{z,\ell}^i} = p_{z,\ell}.$$

From the definition of the actions using the market clearing condition on the asset markets we obtain for the asset prices

- for $\sum_{i=1}^N \bar{\theta}_{z,j}^i = \sum_{i=1}^N \bar{\varphi}_{z,j}^i > 0$

$$\pi_{z,j} = \frac{\sum_{i=1}^N \beta_{z,j}^i}{\sum_{i=1}^N \gamma_{z,j}^i} = \frac{\bar{\pi}_{z,j} \sum_{i=1}^N \bar{\theta}_{z,j}^i}{\sum_{i=1}^N \bar{\varphi}_{z,j}^i} = \bar{\pi}_{z,j}$$

- for $\sum_{i=1}^N \bar{\theta}_{z,j}^i = \sum_{i=1}^N \bar{\varphi}_{z,j}^i = 0$

$$\pi_{z,j} = \frac{\sum_{i=1}^N \beta_{z,j}^i}{\sum_{i=1}^N \gamma_{z,j}^i} = \frac{\bar{\pi}_{z,j} \sum_{i=1}^N \frac{\delta}{N}}{\sum_{i=1}^N \frac{\delta}{N}} = \bar{\pi}_{z,j}.$$

The final allocation of sales and of purchases for asset $j \in \mathcal{J}$ are given by

$$\begin{aligned} \varphi_{z,j}^i &= \gamma_{z,j}^i, \\ \theta_{z,j}^i &= \frac{\beta_{z,j}^i}{\pi_{z,j}}. \end{aligned}$$

The final allocation of good $\ell \in \mathcal{L}$ available for consumption after trading at node $z \in \mathbf{D}$ is given by

$$x_{z,\ell}^i = w_{z,\ell}^i + \sum_{j=1}^J \varphi_{z^-,j}^i C_{j\ell} - q_{z,\ell}^i + \frac{b_{z,\ell}^i}{p_{z,\ell}} - \sum_{j=1}^J \varphi_{z,j}^i C_{j\ell}.$$

Therefore,

$$\varphi_{z,j}^i = \begin{cases} \bar{\varphi}_{z,j}^i & \text{if } \sum_{i=1}^N \bar{\varphi}_{z,j}^i > 0 \\ \frac{\delta}{N} & \text{otherwise} \end{cases}$$

$$\theta_{z,j}^i = \begin{cases} \bar{\theta}_{z,j}^i & \text{if } \sum_{i=1}^N \bar{\theta}_{z,j}^i > 0 \\ \frac{\delta}{N} & \text{otherwise} \end{cases}$$

$$x_{z,\ell}^i = \begin{cases} \bar{x}_{z,\ell}^i & \text{if } \sum_{i=1}^N \bar{\varphi}_{z,j}^i = \sum_{i=1}^N \bar{\theta}_{z,j}^i > 0 \\ \bar{x}_{z,\ell}^i - \sum_{j=1}^J \frac{\delta}{N} C_{j\ell} & \text{otherwise.} \end{cases}$$

It remains to check that the budget constraint ($*_{z^1}^i$) for the bids and offers is satisfied:

$$\sum_{\ell=1}^L b_{z,\ell}^i + \sum_{j=1}^J \beta_{z,j}^i \leq \sum_{\ell=1}^L p_{z,\ell} q_{z,\ell}^i + \sum_{j=1}^J \pi_j \gamma_{z,j}^i + \sum_{j=1}^J \left(\theta_{z^-,j}^i - \varphi_{z^-,j}^i \right) D_{z,j}.$$

Inserting the assumed strategies for $b_{z,\ell}^i$, $q_{z,\ell}^i$, $\gamma_{z,j}^i$ and $\beta_{z,j}^i$ we obtain for ($*_{z^1}^i$)

- for $\sum_{i=1}^N \bar{\theta}_{z,j}^i = \sum_{i=1}^N \bar{\varphi}_{z,j}^i > 0$

$$\begin{aligned} & \sum_{\ell=1}^L \bar{p}_{z,\ell} \left(\bar{x}_{z,\ell}^i + \sum_{j=1}^J \bar{\varphi}_{z,j}^i C_{j\ell} \right) + \sum_{j=1}^J \bar{\pi}_{z,j} \left(\bar{\theta}_{z,j}^i - \bar{\varphi}_{z,j}^i \right) \\ & \leq \sum_{\ell=1}^L \bar{p}_{z,\ell} \left(w_{z,\ell}^i + \sum_{j=1}^J \bar{\varphi}_{z^-,j}^i C_{j\ell} \right) + \sum_{j=1}^J \left(\bar{\theta}_{z^-,j}^i - \bar{\varphi}_{z^-,j}^i \right) D_{z,j} \end{aligned}$$

- for $\sum_{i=1}^N \bar{\theta}_{z,j}^i = \sum_{i=1}^N \bar{\varphi}_{z,j}^i = 0$

$$\sum_{\ell=1}^L \bar{p}_{z,\ell} \bar{x}_{z,\ell}^i \leq \sum_{\ell=1}^L \bar{p}_{z,\ell} \left(w_{z,\ell}^i + \sum_{j=1}^J \bar{\varphi}_{z^-,j}^i C_{j\ell} \right) + \sum_{j=1}^J \left(\bar{\theta}_{z^-,j}^i - \bar{\varphi}_{z^-,j}^i \right) D_{z,j}$$

which holds since $(\bar{x}^i)_{i \in \mathcal{N}}$ was assumed to be a feasible allocation. As $(w_z^i)_i \gg 0$, this strategy profile is full. This completes the proof. \square