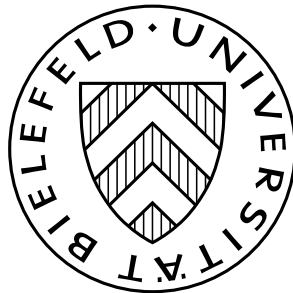# Learning in Infinite Horizon Strategic Market Games with Collateral and Incomplete Information

Sonja Brangewitz and Gaël Giraud

Learning in Infinite Horizon Strategic Market Games with Collateral and Incomplete Information[*]

Sonja Brangewitz[†] and Gaël Giraud[‡]

November 28, 2011

Abstract

We study a strategic market game with finitely many traders, infinite horizon and real assets. To this standard framework (see, e.g. Giraud and Weyers, 2004) we add two key ingredients: First, default is allowed at equilibrium by means of some collateral requirement for financial assets; second, information among players about the structure of uncertainty is incomplete. 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 then 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 sequentially strictly 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

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

Contents

1 Introduction

The events leading to the financial crisis 2007-2008 have highlighted the importance of belief heterogeneity and how financial markets also create opportunities for agents with different beliefs to leverage up and speculate. Several investment and commercial banks invested heavily in mortgage-backed securities, which subsequently suffered large declines in value. At the same time, some hedge funds profited from the securities by short-selling them. One reason for why there has been relatively little attention, in economic theory, paid to heterogeneity of beliefs and how these interact with financial markets is the market selection hypothesis. The hypothesis, originally formulated by Friedman (1953), claims that in the long run, there should be limited differences in beliefs because agents with incorrect beliefs will be taken advantage of, and eventually be driven out of the markets by those with the correct belief. Therefore, agents with incorrect beliefs will have no influence on the economic activity in the long run. This hypothesis has been formalized and extended in recent work by Blume and Easley (2006) and Sandroni (2000). However these authors assume that financial markets are complete, an assumption which plays a central role in allowing agents to pledge all their wealth. By contrast, Cao (2011) presents a dynamic general equilibrium framework in which agents differ in their beliefs but markets are endogenously incomplete because of collateral constraints. Collateral constraints limit the extent to which agents can pledge their future wealth and ensure that agents 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 on the basis of those heterogeneous beliefs. This leads to additional leverage and asset price volatility (relative to a model with homogeneous beliefs or relative to the complete markets economy).

In this paper, we explore a middle ground between these two strands of literature, where traders have heterogeneous beliefs, cannot be simply driven out of the market (thanks to the collateral constraints, as in Cao 2011) but strategically learn the true state of the world. The uncertain state of the world is a transition matrix that gives the probabilities with which a succeeding node in a tree-like time structure is reached. The sets of players and actions are common knowledge, but the distribution of initial endowments and one-period utility levels conditional on action profiles is chosen randomly in each period, and the players do not observe nature's choice. Neither do they observe any player's action —hence, markets are assumed to allow anonymous trading. The probability distribution according to which uncertainty realizes in each period is a (stationary) Markov chain. This Markov distribution itself is chosen at random once and for all at the start of play, and, again, the investors do not observe nature's choice. The players

have a common prior[1] over the finite set of possible Markov chains (states of the world), and they have various ways of learning the state of the world over time. First, each player observes her own initial endowment and realized payoff in each period —both are realizations of random variables whose distribution depends on the state. 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 investors to be able to learn the state, we flesh-out the general equilibrium skeleton with a strategic market game.[2] More precisely, we study a strategic market game with infinite horizon, finitely many long-lived traders, and short-lived real assets. Collateral requirements for financial assets are introduced as in Geanakoplos and Zame (2007) and the subsequent literature. Investors' actions are not observable, so that we stick to the basic anonymity property of large markets. 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 sequentially strictly 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 belief 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.

The double role of financial assets

Our model extends the finite horizon case <u>without</u> default considered in Giraud and Weyers (2004) and the finite horizon <u>with</u> default examined in Brangewitz (2011). In both papers, uncertainty is only on future endowments while, here, we allow for uncertainty on endowments, utilities and asset returns.[3] Moreover, the authors restricted themselves to a very specific game-theoretic set-up: one with partial monitoring (players condition

---

[1]See the footnote on page 21 for an argument of this assumption.
[2]See Giraud (2003) for an introduction.
[3]See Thomas (1995) for an example of general equilibrium model where uncertainty affects consumer's future utilities.

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 considers 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, and the authors proved the analogue of a perfect Folk theorem.

By contrast, we consider perfect Bayesian equilibria where players can update their belief along the play of the game. This deeply changes the strategic challenges at stake: 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 realizations, and updating their beliefs about the region of the event-tree where they are currently located. Financial assets, now, play a double role: On the one hand, they serve as means for reallocating one's resources in face of risky events, on the other they can have a function analogous to that of "arms" in the multiarmed bandit problem (Rothschild, 1974). Since a buyer and a seller of an asset do not know exactly at which node of the tree they are, for each asset, there is a separate unknown probability distribution over returns. 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 return in the short run to gain some information that will help her 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.

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 5 below), market incompleteness does not prevent investors from learning the state. We show, indeed, that, despite price manipulation, infinite-horizon incomplete markets may be fully revealing. This is

in the line with the static general equilibrium literature with real assets, where generically, every equilibrium is fully revealing (Radner 1979, Duffie and Shafer 1985). Beyond the difference between our imperfectly competitive approach and the perfect competition hypothesis, the interpretation of our result, however, strongly differs from that of the literature just mentioned. First, we focus only on fully-revealing equilibria where learning enables players to guess the state in the long-run with an arbitrary accuracy: There might exist plenty other —partially revealing or even non-informative— equilibria. Second, we restrict ourselves to real assets for the sake of clarity. A careful reading of our proof, however, shows that our result goes through in the nominal asset case, as well.[4] Therefore, from the 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). Third, our (partial) Folk theorem implies a huge indeterminacy of the set of strategic equilibria which also contrasts with the generic determinacy obtained by Duffie and Shafer (1985) in the perfectly competitive set-up with incomplete markets of real assets. Fourth, this indeterminacy delivers an ambivalent message in terms of welfare: Many learning equilibria, although they are fully-revealing, are Pareto-dominated by competitive (Radner) equilibria, while many others Pareto-dominate the perfectly competitive benchmark with incomplete markets.[5]

A last point is worth emphasizing before turning to the strategic aspects of our work. 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 set-up, there is no need for such any extra 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 do not to lend her money so many times — hence, participating to 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.

---

[4]The proof is actually even simpler. This is why we have treated the real asset case.

[5]As 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 micro-structure and to allow players to send limit-price (not just market) orders to the clearing house.

Asymmetric information and markets

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 our approach (the Informativeness Assumption, (IA)) can be stated as follows: Observing the realization of one's (random) initial endowments, one-period (random) utility levels and (strategically determined) final allocations together with *all* the assets' returns suffices for every single trader to learn the true state of the world in the long-run with probability arbitrarily close to 1. 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 1 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 assets' returns in certain periods, then she might not draw the right conclusion about which Markov chain is driving uncertainty, so that 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 undistinguishable 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-set-up, 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 related to, for example, job market signaling model of Spence (1973) or the competitive insurance market considered in Rothschild and Stiglitz (1976). However, we do not consider a classical principal agent model. Every individual may act as a seller or a buyer (or both simultaneously), and this on commodity as well as asset markets. Therefore, we cannot impose, for example, that a seller is always less informed than a buyer or vice versa. Finally, we consider only finitely many players. Our set-up 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 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 first (partial) Folk theorem under the simplifying assumption of complete information. Section 5 extends the later result to the incomplete information case. The last section concludes.

2 The Markov Strategic Market Game with Collateral

2.1 The Markov Economy

The environment

Uncertainty about future states is modeled in a Markov set up, following Cao (2011). We assume that in each period, t, the state of nature in the next period is chosen using a Markov transition matrix with a finite set of possible states of nature $\mathcal{S} = \{1, ..., S\}$. Therefore, the state tomorrow only depends on the state today and not the whole history of states that were realized in the past. Nevertheless as in Magill and Quinzii (1994) and the subsequent literature, time, uncertainty and the revelation of information can be described by an event tree, i.e., a directed graph $(\mathbf{D}, \mathcal{A})$ consisting of a set $\mathbf{D}$ of vertices and a set $\mathcal{A} \subset \mathbf{D} \times \mathbf{D}$ of (oriented) arcs.[6] In our Markov set-up, we assume that each node ξ has the same outdegree $S > 1$, and the choice of nodes adjacent from ξ is governed by a Markov chain. A node ξ can be interpreted as a date-event pair (t, s_{t-1}, s), where $t \geq 1$ is the minimal length of a walk between ξ_0 and ξ, $s_{t-1} \in \prod_{t'=1}^{t-1} \mathcal{S}$ 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 $\tau(\xi)$ be the time at which node ξ is reached, i.e. $\tau : \mathbf{D} \to \mathbb{N}$ such that $\xi = (t, s_{t-1}, s) \mapsto t$. Define a partial order $\geq$ on $\mathbf{D}$ by $\xi = (t, s_{t-1}, s) \geq \xi' = (t', s_{t'-1}, s')$ if, and only if, there is a walk from ξ' to ξ. Of course, if $\xi \neq \xi'$ and $\xi \geq \xi'$, then $\xi > \xi'$. The unique predecessor of ξ is denoted by $\xi^- = (t-1, s_{t-2}, s')$.[7] The set of immediate successors of ξ, denoted by ξ^+, is the set of nodes that are adjacent from ξ. For any node $\xi \in \mathbf{D}$, the set of all nodes with $\xi' \geq (>)\xi$ is denoted by $\mathbf{D}(\xi)$ $(\mathbf{D}(\xi)^+)$ and is itself a tree with root ξ.

A state of the *world* corresponds to a transition matrix, ω, 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

We consider a pure exchange economy $\mathcal{E}$ with a finite set, $\mathcal{N} = \{1, ..., N\}$, of individuals, L consumption goods, usually indexed by ℓ, and J short-term real assets, indexed by j. The

---

[6] The vertex (or node) ξ can be thought of as a particular state of nature and time. If (ξ, η) is an arc, η is a node that directly follows ξ. Formally, ξ is adjacent *to* η and η is adjacent *from* ξ. The number of nodes adjacent to a given vertex ξ is the *indegree* of ξ, i.e. the number of immediate (or direct) predecessors; the number of nodes adjacent from ξ, its *outdegree*, i.e. the number of direct followers. A *walk* from ξ_1 to ξ_k is a sequence $(\xi_1, \xi_2, ..., \xi_k)$ in $\mathbf{D}$ such that ξ_i is adjacent to ξ_{i+1} for $1 \leq i \leq k-1$. There is a unique root ξ_0 (whose indegree is zero). Each node, except the root, has indegree equal to 1, and there is no cycle in $\mathbf{D}$.

[7] We define $s_{-1} = \emptyset$.

possibility of default is introduced by a collateral requirement as in Araujo et al. (2002). A financial asset $j \in \mathcal{J} := \{1, ..., J\}$ is characterized by a tuple (ξ^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 ξ^j. The promised amount of goods is described by a function $A_j : \mathbf{D} \to \mathbb{R}_+^L$ such that $A_j(\xi) = 0$ for all $\xi \in \mathbf{D} \setminus \left(\xi^j\right)^+$. For $\xi' \in \left(\xi^j\right)^+$, the promises $A_j(\xi')$ are the amounts 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 ξ^j. The delivery, $p_\xi \cdot A_j(\xi)$, is assumed to be made in fiat money using spot prices, $p_\xi \in \mathbb{R}_+^L$. We only consider short-term assets. Therefore, for other nodes before the issuing node and at least two periods after the asset was issued, we assume that the promised amounts are zero. The vector $C_j \in \mathbb{R}_+^L$ is the amount of collateral needed at the issuing node, ξ^j, in order to back up the promised delivery A_j. Only consumption goods can serve as collateral.[8] Commodities are assumed to be **perishable**. Thus, they have to be consumed at the very date they enter the economy (as initial endowment), unless they are stored as collateral. Individuals are not allowed to consume a collateral, which is stored in a warehouse for one period. For simplicity, after having been stored one period, a collateral must be consumed, otherwise it gets lost.[9] For our Markov environment, we assume that at each node $\xi \in \mathbf{D}$ the "same" finite number of financial assets is issued. As the time horizon is infinite there will be infinitely many assets in total.

The players

Every player $i \in \mathcal{N}$ is characterized by a twice continuously differentiable, strictly increasing and concave utility function $u_\xi^i : \mathbb{R}_+^L \to \mathbb{R}$ and a strictly positive initial endowment in consumption goods $w_\xi^i \in \mathbb{R}_{++}^L$ at every node $\xi \in \mathbf{D}$. We assume that $\left(u_\xi^i(\cdot)\right)_\xi$ are uniformly bounded below for all individuals i. Therefore, without loss of generality suppose $u_\xi^i(0) = 0$. Moreover, we assume that individual endowments are uniformly bounded above by some $\overline{w}$, across individuals and periods. Initial holdings of assets are 0. Player i maximizes her expected, discounted utility from consumption. This expectation depends on her subjective beliefs on the state of the world $\omega \in \Omega$, which may themselves vary across time, depending upon the signals sent by other players during the play of the game. We shall therefore define player's i objective function after having recalled the basic structure of the strategic market game.

We also denote by $\mathcal{E}_\xi = \langle w_\xi^i, u_\xi^i(\cdot), (\xi^j, A_j, C_j)_{j \,|\, \xi^j = \xi} \rangle$ the finite-dimensional *one-shot economy* at node ξ. We denote the infinite horizon economy starting from a certain node ξ, that is not necessarily the root ξ_0, for short *the economy after ξ*, by $\bigcup_{\xi' > \xi} \mathcal{E}_{\xi'}$.

---

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

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

2.2 The Strategic Market Game with Collateral

At each period, players take part to a strategic market game *à la* Shapley and Shubik (1977): Each individual places for every consumption good $\ell \in \mathcal{L}$ at every node $\xi \in \mathbf{D}$ a bid $b^i_{\xi,\ell}$ and an offer $q^i_{\xi,\ell}$. The bid $b^i_{\xi,\ell}$ signals how much (in terms of fiat money) player i is willing to pay for the purchase of good ℓ and the offer $q^i_{\xi,\ell}$ (in terms of physical commodities) is the amount she wants to sell. The price of good ℓ is then computed as the ratio of the total bid to the total offer, that is

$$
p_{\xi,\ell} = \begin{cases} \frac{\sum_{i=1}^{N} b^i_{\xi,\ell}}{\sum_{i=1}^{N} q^i_{\xi,\ell}} & \text{if } \sum_{i=1}^{N} q^i_{\xi,\ell} > 0 \\ 0 & \text{otherwise} \end{cases}
$$

A market without trade is said to be closed.[10]

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

$$
\pi_{\xi,j} = \begin{cases} \frac{\sum_{i=1}^{N} \beta^i_{\xi,j}}{\sum_{i=1}^{N} \gamma^i_{\xi,j}} & \text{if } \sum_{i=1}^{N} \gamma^i_j(\xi) > 0 \\ 0 & \text{otherwise} \end{cases}
$$

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:

$$
D_{\xi',j} = \min \{ p_{\xi'} \cdot A_j(\xi'), p_{\xi'} \cdot C_j \} \tag{D}
$$

at node $\xi' \in (\xi^j)^+$. Hence, whether default appears or not is not the outcome of a strategic decision but depends upon the commodity price $p_{\xi'}$, which is strategically determined by bids and offers posted at node $\xi' \in (\xi^j)^+$.

Feasible bids and offers

Some physical and budgetary restrictions are put on the bids and offers individuals can choose. At every node $\xi \in \mathbf{D}$ and for every financial asset, player i needs to own the required amount of collateral, which depends on the quantity of asset offered for sales and *not* on the net trades.[11] Assuming player i offers to sell $\gamma^i_{\xi,j}$ units of asset j at node

---

[10]Defining the price as zero when there are no offers on the market we follow here for example Amir et al. (1990, p.128). Similar assumptions can be found in Postlewaite and Schmeidler (1978, p.128), Peck et al. (1992, p.275) or Giraud and Weyers (2004, p.474).

[11]As discussed in Dubey and Geanakoplos (2003), netting before imposing the collateral requirement would suppress any constraint on the size of short sales. This would make the proof of our partial Folk

ξ, then she needs to store $\gamma^i_{\xi,j} C_j \in \mathbb{R}^L_+$ as collateral.[12]

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

$$\sum_{j=1}^J \gamma^i_{\xi,j} C_{j\ell} \leq w^i_{\xi,\ell} \qquad (F1\xi)$$

and

$$q^i_{\xi,\ell} \leq \sum_{j=1}^J \gamma^i_{\xi^-,j} C_{j\ell} + \Delta(F1\xi), \qquad (F2\xi)$$

where $\Delta(F1\xi)$ stands for the difference between the right-hand side and the left-hand side of $(F1\xi)$. Inequality $(F1\xi)$ says that the collateral that can be stored by i at node ξ must be taken out of initial endowments. In particular, it cannot consist of commodities that are already inherited from the past as collaterals. This is a way to capture our assumption that every collateral lives at most one period. Either it is consumed at the 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 $(F2\xi)$ 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 $\xi \in \mathbf{D}$ plus the collateral that was put aside in the previous period. Of course, we impose:

$$q^i_{\xi,\ell}, b^i_{\xi,\ell}, \beta^i_{\xi,j}, \gamma^i_{\xi,j} \geq 0 \qquad (F3\xi)$$

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:

$$\sum_{\ell=1}^L b^i_{\zeta,\ell} + \sum_{j=1}^J \beta^i_{\zeta,j}$$

---

theorem only easier.

[12]Later, on page 13 when defining the final allocation in consumption goods, the collateral requirement is taken using the final asset sales, denoted by $\varphi^i_{\xi,j}$ and not directly on the offers $\gamma^i_{\xi,j}$.

$$\leq \sum_{\ell=1}^{L} p_{\zeta,\ell} q_{\zeta,\ell}^{i} + \sum_{j=1}^{J} \pi_{\zeta,j} \gamma_{\zeta,j}^{i} + \sum_{j=1}^{J} \left(\theta_{\zeta^{-},j}^{i} - \varphi_{\zeta^{-},j}^{i} \right) D_{\zeta,j} \qquad (*_{\xi}^{i}1)$$

for all $\zeta \leq \xi$ where $\theta_{\zeta^{-},j}^{i}$ denotes the final asset purchases and $\varphi_{\zeta^{-},j}^{i}$ the asset sales at node ζ^{-} (as it will be defined below). Thus, by condition $(*_{\xi}^{i}1)$ 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, $\theta_{\zeta^{-},j}^{i} - \varphi_{\zeta^{-},j}^{i}$. As soon as $(*_{\xi}^{i}1)$ is violated, say at node ξ, individual i is removed from the game for all subsequent nodes $\mathbf{D}^{+}(\xi)$, and all her goods are confiscated forever.

We shall also need the following condition, for every i:

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

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

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

and her purchases

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

Note that if $\theta_{\xi,j}^{i} - \varphi_{\xi,j}^{i} < 0$ then player i sold more of the financial asset $j \in \mathcal{J}$ than she bought. Analogously for $\theta_{\xi,j}^{i} - \varphi_{\xi,j}^{i} > 0$ she is a net buyer.

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

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

Remark 1. To the best of our knowledge, condition $(*_{\xi}^{i}2)$ is new in the strategic market game literature. It seems to us natural once collateral requirements are introduced.

Suppose, indeed, that individual i is the only one who wants to trade on the financial markets, i.e., $\sum_{k \neq i} \gamma_{\xi,j}^k = \sum_{k \neq i} \beta_{\xi,j}^k = 0$. Absent condition $(*_\xi 2)$, 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 next period just by trading "with herself" today. If for several periods such a strategy is played, while the other players play zero strategies, this would conflict with our assumption that commodities are perishable.

3 Feasibility and *interim* individual rationality

Allowable strategies

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

$$A_\xi^i = \left\{ \left(q_{\xi,\ell}^i, b_{\xi,\ell}^i\right)_{\ell \in \mathcal{L}}, \left(\gamma_{\xi,j}^i, \beta_{\xi,j}^i\right)_{j \in \mathcal{J}} \in \mathbb{R}_+^{2L} \times \mathbb{R}_+^{2J} \big| (F1\xi), (F2\xi) \text{ and } (F3\xi) \text{ are satisfied} \right\}.$$

Notice that A_ξ^i depends upon ξ but not upon ω. Let $A_\xi := \prod_{i=1}^N A_\xi^i$. Note that the definition of an action set includes actions that possibly violate the budget constraint $(*_\xi^i 1)$ or $(*_\xi^i 2)$.[13] The stage-payoff of player i at node $\xi = (t, s_{t-1}, s)$ is given by the utility, $u_\xi^i(x_\xi^i)$, she obtains from consumption.

Prices are publicly observed by every player. The information transmitted through prices is therefore common knowledge. However, at each node, every player also observes her own initial endowment, her final stage-payoff, her final allocation as well as the returns of the assets present in her portfolio. These observations constitute the *private history* of player i. A *strategy* of player i consists in choosing an action at every node $\xi \in \mathbf{D}$ as a function of her own private history. Let H_ξ^i denote the set of possible private histories for individual i at node ξ, given by

$$H_\xi^i := \left\{ \left(p_{\xi'}, \pi_{\xi'}, \varphi_{\xi'}^i, \theta_{\xi'}^i, x_{\xi'}^i, u_{\xi'}^i(x_{\xi'}^i), w_{\xi'}^i, w_\xi^i\right) \big| \forall \xi' < \xi \right\}.$$

The history at the root ξ_0 is given by $H_{\xi_0}^i = \left\{ w_{\xi_0}^i \right\}$. Formally, a strategy of player i is a map

$$\sigma^i : \bigcup_{\xi \in \mathbf{D}} H_\xi^i \to \left(\mathbb{R}_+^L\right)^2 \times \left(\mathbb{R}_+^J\right)^2$$

such that $\sigma^i(h) \in A_\xi^i$ for all $\xi \in \mathbf{D}$ and for all $h \in H_\xi^i$. Actions are not observed along the play of the game, which contrasts with the setting considered, e.g., by Wiseman (2011).

---

[13]An alternative would consist in incorporating these constraints into the very definition of a player's strategy set but this would lead to a generalized game as introduced by Debreu (1952) (see also Harker (1991) or Facchinei and Kanzow (2010)).

Remark 2. As is well-known, strategic market games exhibit no-trade as a one-shot Nash equilibrium.[14] As we want to prove the analogue 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 that would allow to get rid of the autarkic one-shot equilibrium[15]), 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. 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.

Definition 3.1 (**Full strategy profile**). A strategy profile $\sigma := (\sigma^i)_i$ is called *full* if, the following holds

$$\sum_{i=1}^{N} q^i_{\xi,\ell} > 0, \quad \sum_{i=1}^{N} b^i_{\xi,\ell} > 0, \quad \sum_{i=1}^{N} \gamma^i_{\xi,j} > 0, \quad \sum_{i=1}^{N} \beta^i_{\xi,j} > 0$$

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

Private *interim* beliefs

At each node ξ, payoffs are determined as follows: action profile $a_\xi \in A_\xi$ is played; it induces, say, x^i_ξ as a final allocation for player i —which is observed by i only. Then player i's random payoff, $u^i_\xi(x^i_\xi)$, which is also observed by player i only, is drawn according to ω. Notice that, when entering at node ξ, player i may not know for sure that the current node is ξ. 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, assets' returns. So player i's belief about player j's private payoff and other higher-order beliefs are unrestricted. Let $\mathbb{P}^i_\xi(h^i_\xi) \in \Delta(\Omega)$

---

[14]See Weyers (2004) for the elimination of this autarkic equilibrium after two rounds of elimination of dominated strategies.

[15]Such 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.

denote player i's *private belief* at node ξ.[16] Together with a strategy profile, σ, such a probability $\mathbb{P}_\xi^i(h_\xi^i)$ induces a distribution $\mathbf{P}_\xi^i(h_\xi^i, \sigma)$ (or $\mathbf{P}_\xi^i(\sigma)$ in short) over the random characteristics of the economy to be selected after ξ, i.e., over $\bigcup_{\xi' > \xi} \mathcal{E}_{\xi'}$. In particular, it provides a distribution over i's future payoffs which, by a slight abuse of notations, is also denoted $\mathbf{P}_\xi^i(\sigma)$. At each node, whatever being the past history, individuals are supposed to maximize her expected, discounted utility using their private *interim* belief and a common discount factor $\lambda \in [0, 1]$.[17] The objective function of player i is therefore of the form

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

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

Feasible allocations and *interim* individual rationality

Without considering explicitly actions or strategies we define feasible allocation as follows:

Definition 3.2 (**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 conditions are satisfied:

- *Individual budget restriction for every player i and every node $\xi \in \mathbf{D}$:*[18]

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

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

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

- *market clearing on financial markets for every asset $j \in \mathcal{J}$ and every node:*
$$
\sum_{i=1}^N \bar{\theta}_{\xi,j}^i = \sum_{i=1}^N \bar{\varphi}_{\xi,j}^i
$$

---

[16]Hereby, $\Delta(\Omega)$ is the set of all probability distributions over the finite set of states of the world.
[17]Allowing for idiosyncratic discount factors would only require notational changes.
[18]We define $\bar{\varphi}_{\xi_0^-,j}^i = \bar{\theta}_{\xi_0^-,j}^i = 0$.

16

- *and feasible trade in financial assets for every good $\ell \in \mathcal{L}$, every node and every player i:*
  $\sum_{j=1}^{J} \bar{\varphi}_{\xi,j}^{i} C_{j\ell} \leq w_{\xi,\ell}^{i}$

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, $u_{\xi}^{i}(\cdot)$, 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 $\bar{u} := \max_{i \in \mathcal{N}} \bar{u}^{i}$.

In the next definition, individual rationality is understood according to the *interim* private beliefs shared by players along the play of the game. It is therefore defined *given* some state of the world, ω, and some strategy profile, σ.

Definition 3.3 (Sequentially strictly individually rational allocation).

A feasible allocation $(\bar{x}^{i})_{i \in \mathcal{N}}$ is said to be *sequentially strictly individually rational (*SSIR*)* *given* ω, if

$$U_{\mathbf{D}(\xi)}^{i}(x^{i}, \sigma, \omega) > (1 - \lambda) E_{\mathbf{P}_{\xi}^{i}(\sigma)} \sum_{\xi' = (t, s_{t-1}, s) > \xi} \lambda^{t-1} u_{\xi'}^{i}(w_{\xi'}^{i}).$$

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 1. *If the initial allocations $(w_{\xi}^{i})_{i} \gg 0$ are Pareto-inefficient in the L-good spot economy at each node $\xi \in \mathbf{D}$, then the economy $\mathcal{E}$ admits a sequentially strictly individually rational and feasible (*SSIRF*, for short) allocation.*

The next Lemma will prove useful for our main result. It shows that every SSIRF allocation can be enforced by means of some adequate strategy. Such a strategy, however, need not fulfill any equilibrium requirement.

Lemma 2. *Let $(\bar{x}^{i})_{i \in \mathcal{N}}$ be a *SSIRF* 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 utility of this strategy profile is arbitrarily close to the node-wise utility of $(\bar{x}^{i})_{i \in \mathcal{N}}$. Whatever being the past history, play*

17

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

$$q_{\xi,\ell}^i = w_{\xi,\ell}^i + \sum_{j=1}^J \bar{\varphi}_{\xi^-,j}^i C_{j\ell}$$

$$b_{\xi,\ell}^i = \bar{p}_{\xi,\ell} \left(\bar{x}_{\xi,\ell}^i + \sum_{j=1}^J \bar{\varphi}_{\xi,j}^i C_{j\ell} \right)$$

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

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

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

If we target a given allocation using the full strategies as defined in Lemma 2 and this allocation does not always require trade on the asset markets, then we cannot target the allocation exactly. For the details we refer to the proof in 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 close to the target allocation.

4   Complete Information

We first state our result in the simpler case where information is complete, i.e., the Markov chain ω is known from the beginning by every player.

Theorem 1. *Suppose that $\#\Omega = 1$. Every allocation that is* SSIRF *can be approximately enforced as a subgame perfect Nash equilibrium (SPNE).*

Proof. Let $(x_\xi^{*i})_{i,\xi}$ be a SSIRF allocation for the transition matrix ω with stage-payoffs $(v_\xi^{*i})_{i,\xi} := \left(u_\xi^i(x_\xi^{*i}) \right)_{i,\xi}$. We denote by E_ω the expectation operator with respect to the beliefs that the state of the world is given by ω. The utility for player i resulting from x^{*i} is then given by

$$U_{\mathbf{D}(\xi_0)}^i(x^{*i}, \sigma, \omega) = (1-\lambda) \sum_{\xi'=(t,s_{t-1},s) > \xi_0} \lambda^{t-1} E_\omega \left[u_{\xi'}^i(x_{\xi'}^{*i}) \right].$$

We construct a sequence of payoff vectors $\left((v^{i,n\text{dev}})_i \right)_{n \in \mathbb{N}}$ that result from SSIRF allocations, and such that: $v_\xi^{i,(n+1)\text{dev}} < v_\xi^{i,n\text{dev}}$ for every integer $n \in \mathbb{N}$ and every node ξ —

with $v^{i,0\text{dev}} = v^{*i}$ for each i. These payoffs will be the long-run payoffs after n deviations. They are constructed as follows:

$$v_\xi^{i,n\text{dev}} := u_\xi^i \left(x_\xi^{i,n\text{dev}} \right) \text{ with } x_\xi^{i,n\text{dev}} := \rho_n x_\xi^{*i} + (1 - \rho_n) w_\xi^i, \quad \rho_n \in (0,1).$$

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

$$0 < \varepsilon_n < v_\xi^{i,n\text{dev}} - v_\xi^{i,(n+1)\text{dev}} \tag{1}$$

Using Lemma 2 we construct full strategies that result approximately in the target allocation $(x^{*i})_{i \in \mathcal{N}}$. If there is no deviation from these strategies, then every individual continues to play these strategies. The punishment, if one individual deviates, is to play the following strategies: Every individual bids and offers $\frac{\delta}{N}$ with $\delta > 0$ small on the goods and on the assets markets for the next T_n periods, if the nth deviation had been observed. As all individuals bid and offer the same quantities, these strategies mimic the no trade equilibrium and everybody keeps her initial endowment. On the asset markets however every individual sells $\frac{\delta}{N}$ of every asset and hence needs to have a collateral of $\frac{\delta}{N} C_{j\ell}$. As there is no trade on the goods markets, this additional collateral needs to be established from the initial endowments, which are strictly positive. Thus, δ needs to be small enough such that this is can be done.

After the punishment phase dedicated to the nth deviation there is a reward phase, if no further deviation has occurred. As soon as another deviation occurs, a new punishment phase of length T_{n+1} starts immediately. Suppose the nth deviation has occurred and there was no further deviation during the punishment phase. Then in the reward phase the individuals play some actions, as defined in Lemma 2, leading approximately to a SSIRF allocation with a stage payoff of $v_\xi^{i,n\text{dev}}$. Notice that in order to settle the asset market obligations from the punishment phase and to establish the right asset holdings to reach $v_\xi^{i,n\text{dev}}$ two periods of transition are required to ensure that the individual budget constraint $(*_\xi^i 1)$ is not violated. For the details concerning the transition periods we refer to the proof of Theorem 2, page 31. Taking this punishment behavior into consideration we show that there is no incentive to deviate.

Suppose individual i deviates at node $\xi = (t', s_{t'-1}, s)$ and this was the $(n+1)$th deviation observed. We need to compare the gains and losses from the deviation. Individual i can by deviating maximally reach the upper bound of her utility given by $\bar{u}^i$ in the period of her deviation. In the succeeding T_{n+1} periods after the deviation: According to the definition of the strategies above she stays close her initial endowment. The $(n+1)$th

deviation payoff is arbitrarily close to

$$(1 - \lambda)\Big[\lambda^{t'-1}\overline{u}^i + \sum_{t=t'+1}^{t'+T_{n+1}} \lambda^{t-1} E_\omega[u_\xi^i(w_\xi^i)]$$

$$+ \sum_{t=T_{n+1}+t'+1}^{T_{n+1}+t'+2} \lambda^{t-1}\overline{u}$$

$$+ \sum_{t\geq T_{n+1}+t'+3} \lambda^{t-1} E_\omega[v_\xi^{i,(n+1)\mathrm{dev}}]\Big]. \tag{2}$$

The long-run discounted payoff after the $(n+1)$th deviation consists of once a (maybe) very high payoff from deviating, then the payoff from a punishment phase lasting T_{n+1} periods, two periods of transition with a payoff of maximally $\overline{u}$ and finally the $(n+1)$th reward payoff.

By contrast, if the $(n+1)$th deviation did <u>not</u> take place, i's long-run payoff starting at time t' would be arbitrarily close to:

$$(1 - \lambda)\Big[\sum_{t\geq t'} \lambda^{t-1} E_\omega[v_\xi^{i,n\mathrm{dev}}]\Big]. \tag{3}$$

Therefore to show that (3) - (2) is positive it is enough to ensure that:

$$\underline{u} - 3\overline{u} + \sum_{t=t'+1}^{t'+T_{n+1}} \lambda^{t-t'-1} E_\omega[v_\xi^{i,n\mathrm{dev}} - u_\xi^i(w_\xi^i)] + \varepsilon_n\Big[\sum_{t\geq T_{n+1}+t'+3} \lambda^{t-t'-1}\Big] > 0.$$

Note that since $v^{i,n\mathrm{dev}}$ was assumed to be a payoff that results from a sequentially strictly individually rational allocation we have $E_\omega[v_\xi^{i,n\mathrm{dev}} - u_\xi^i(w_\xi^i)] > 0$ for every $t \in \mathbb{N}$, for every individual $i \in \mathcal{N}$. Therefore define

$$g_\xi := \min_{i\in\mathcal{N}} E_\omega[v_\xi^{i,n\mathrm{dev}} - u_\xi^i(w_\xi^i)]$$

Therefore, it is sufficient to require that:

$$\underline{u} - 3\overline{u} + \sum_{t=t'+1}^{t'+T_{n+1}} \lambda^{t-t'-1} g_\xi + \varepsilon_n \frac{\lambda^{T_{n+1}+2}}{1-\lambda} > 0.$$

It is easy to see that, whatever being the distance, $\underline{u} - 3\overline{u}$, and for every $\varepsilon_n > 0$, there exists some T_{n+1} big enough so that this last inequality is satisfied. Hence, deviating behavior is not profitable. This completes the proof. □

5 Incomplete Information

In this section, we turn to the general case where $\#\Omega \geq 1$. Players observe neither the choice of ω, nor that of ξ. They 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.[19] Each household has *five* ways of updating its beliefs about the state ω over time.

- First, at node ξ, each player privately observes her own (random) spot endowment, w_ξ^i, which is chosen by nature according to the transition matrix ω.

- Second and third, at every node, after having played her action, each player observes public prices, p_ξ and π_ξ, together with her final allocation, x_ξ^i. Prices and final allocations depend upon the players' actions and vary in informativeness across action profiles: they only reveal the part of the privately hold information that players are ready to transmit through their bids and offers.

- Fourth, a trader may also learn about the state by observing her final stage-payoff, $u_\xi^i(x_\xi^i)$, which is selected according to ξ —given x_ξ^i.

- Finally, the return of the assets she owns in her portfolio (either as a creditor or as a debtor) also provide information about the realization of ξ, hence, about ω.

In order to cope with this differential information set-up, we shall need two key restrictions —Assumptions **G** and **IA**.

> **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. We use this partition of the commodities to ensure that, during the play of the game, a single player cannot prevent the other individuals from learning the true state of the world, ω.

Along a play of the game, while endowments, utility payoffs and asset payoffs are observed privately, prices are publicly revealed. Notice that, given actions a_ξ, prices are entirely determined — i.e., there is no additional randomness on public signals, by contrast with Wiseman (2011) where public signals are random. Of course, the distribution

---

[19]This is in accordance with the arguments provided by Heifetz (2006) showing that it makes hardly 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.

matrix, ω, might be degenerate so that payoffs or returns are non-stochastic conditional on ω. In this case, the realization of payoffs and/or returns perfectly reveals the state of the world. On the contrary, if two distributions do have the same support, players may never be able to learn the true state for sure by just observing their private characteristics and the assets' returns.

Given the state of the world, ω, a strategy profile, $\sigma = (\sigma^i)$, induces a unique probability distribution on the space of sequences $(u_\xi^i(x_\xi^i(\sigma)), x_\xi^i(\sigma), w_\xi^i, A_j(\xi))_{i,j,\xi}$. 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 $(u_\xi^i(x_\xi^i(\sigma)), x_\xi^i(\sigma), w_\xi^i, A_j(\xi))_{i,j,\xi}$ 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 stage-payoff profiles, $(u_\xi^i(\cdot))_{i,\xi}$, final allocations, $(x_\xi^i(\sigma))_{i,\xi}$, endowments, $(w_\xi^i)_{i,\xi}$, and asset returns, $(A_j(\xi))_{j,\xi}$, 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 being the strategy played. Neither need prices suffice to identify by themselves the state.[20] The following assumption is therefore, admittedly, a restriction: it says that, for every "reasonable" strategy profile, stage-payoffs, final allocations and asset returns plus individual endowments contain all the relevant information about ω. Illustrations of textbook models that satisfy this assumption were given in the Introduction of the paper.

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

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

Suppose that the Markov chain ω is irreducible and aperiodic.[21] Then, it admits an invariant measure if, and only if, every state of nature $s \in \mathcal{S}$ is positive recurrent.[22] In this case, μ_ω is unique.

---

[20]When prices are interpreted as public signals, this generality contrasts with Wiseman (2005) where the sole observation of public signals suffices to identify the state with no ambiguity.

[21]A 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 aperiodic. If every state $s \in \mathcal{S}$ is aperiodic, ω is said *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.

[22]A state s is recurrent if, given that the chain starts in s, it will return to s in finite time with probability 1. s is *positive recurrent* if, in addition, the expectation of this hitting time is finite.

Informativeness Assumption (IA)

(1) For any pair of nodes $(t, s_{t-1}, s) = \xi \neq \xi' = (t, s_{t-1}, s')$, any player i, and any strategy profile, σ, that induces an SSIRF allocation at both states, the vectors of signals, $(u_\xi^i(x_\xi^i(\sigma)), x_\xi^i(\sigma), w_\xi^i, A_j(\xi))$ and $(u_{\xi'}^i(x_{\xi'}^i(\sigma)), x_{\xi'}^i(\sigma), w_{\xi'}^i, A_j(\xi'))$ 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 $\mu_\omega = \mu_{\omega'} \Rightarrow \omega \neq \omega'$.

(IA-1) says that, for a reasonable strategy profile, at the end of each period t, each player knows for sure at which node, $\xi = (t, s_{t-1}, s)$, she was playing. 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 on the long-run should be identified. The last section of the paper provides some hints about how this assumption can be weakened.

Definition 5.1 (Perfect Bayesian equilibrium). A pair $\left((\sigma)_{i \in \mathcal{N}}, \left(\mathbb{P}_\xi^i(h_\xi^i) \right)_{i \in \mathcal{N}} \right)$ consisting of a feasible allocation and a system of private beliefs is a perfect Bayesian equilibrium (PBE) if

- $(\sigma)_{i \in \mathcal{N}}$ is sequentially rational given the private beliefs $\left(\mathbb{P}_\xi^i(h_\xi^i) \right)_{i \in \mathcal{N}}$, i.e., starting at any arbitrary node, given the continuation strategies of the other individuals, no individual can improve her utility by unilaterally changing her strategy profile given her private beliefs $\mathbb{P}_\xi^i(h_\xi^i)$,

- and the private beliefs $\left(\mathbb{P}_\xi^i(h_\xi^i) \right)_{i \in \mathcal{N}}$ are updated via Bayes rule whenever it is possible.[23]

Our main result is that, for any strategy profile that yields an allocation of commodities, assets and collaterals that is SSIRF, 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 households are patient enough. Moreover, along such an equilibrium path, every player learns the realized state with arbitrary precision.

---

[23]Due to our assumptions on the Markov chain, every state of nature is reached with a strictly positive 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.

Theorem 2. *Under (G) and (IA), let $\varepsilon > 0$ and $(x^{*i}[\omega])_{i \in \mathcal{N}, \omega \in \Omega}$ be a SSIRF 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 $\lambda(\mathbb{P}) < 1$ 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 $\left(U^i_{\mathbf{D}(\xi_0)}(x^{*i}, \sigma, \omega) \right)_i$. In equilibrium, conditional on ω, each player i's interim private belief converges to the truth: $\lim_{t \to \infty} \mathbb{P}^i_{\xi=(t, s_{t-1}, s)}(h^i_\xi)[\omega] = 1$ with probability 1.*

Proof of Theorem 2

Outline of the proof

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

The equilibrium path uses "blocks" of $M + T$ periods each. An *equilibrium block* has a "target allocation" in commodities, denoted by

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

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

$$\left((\varphi^i_\xi[\omega], \theta^i_\xi[\omega])_i \right)_{\tau(\xi)=1,\dots,M+T}.$$

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*, $\mathbb{P}^i_{\xi \setminus \bar{\xi}}(h^i_\xi) \in \Delta(\Omega)$, defined as follows: Suppose that the block under scrutiny started at node $\bar{\xi} \leq \xi$. Given some history, h^i_ξ, consider the truncated history, $h^i_{\xi \setminus \bar{\xi}}$, containing only the information delivered from $\bar{\xi}$ to ξ^-. The truncated belief, $\mathbb{P}^i_{\xi \setminus \bar{\xi}}\left(h^i_\xi \right) \in \Delta(\Omega)$, is the resulting updated belief starting with prior $\mathbb{P}$ at node $\bar{\xi}$.

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 $\mathbb{P}^i_{\xi \setminus \bar{\xi}}\left(h^i_\xi \right)$, and, in the remaining T periods, households choose a full action profile that yields a stream of final allocations close to the target,

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

with utility payoffs close to

$$\left(u_\xi^i \left(x_\xi^{*i}[\hat\omega] \right)_i \right)_{\tau(\xi)=M+1,\ldots,M+T} .$$

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. 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 modify 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 $\xi = (t, s_{t-1}, s)$, is made arbitrarily close to the initial endowment, w_ξ^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 decreasing in the number of deviations so that $u_\xi^i \left(x_\xi^{i,n\mathrm{dev}}[\omega] \right) <$ $u_\xi^i \left(x_\xi^{i,(n-1)\mathrm{dev}}[\omega] \right)$ for each node $\xi = (t, s_{t-1}, s)$ of the post-deviation block and each state of the world, ω —where n is the number of deviations already observed.[24] That is, the payoff to a deviator is lower than she would get in equilibrium, regardless of the

---

[24]Of course, $u_\xi^i \left(x_\xi^{i,0\mathrm{dev}}[\omega] \right) = u_\xi^i \left(x_\xi^{*i}[\omega] \right)$.

state. A patient player, therefore, will not deviate, neither on, nor off the equilibrium path, regardless of her private beliefs.

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 A; but player's 1 private information at node ξ indicates state B more strongly. Player 1 believes that eventually everybody's belief will converge to a Dirac mass on state 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 that no further learning occurs. This happens, for example, if, from ξ on, individual endowments no more depend upon the state, $\xi' > \xi$, 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 B that the expected time before convergence is very long, even whenever the equilibrium path does call for further experimentation.

Further, in state B, the equilibrium actions specified for state A may yield a lower stage-payoff to player 1 than her initial endowments in state B, i.e., than the actions designed to punish player 1 for a deviation in state 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 B, so that the market belief may adjust toward state B. Then, such a deviation may be profitable for player 1 *even when her private information is consistent with state A*, provided the punishment profile specified in state B gives her a higher payoff when the actual state is A than does the on-path profile specified in state A. This can occur, again, if player 1's post-deviation payoff in B is higher than the final allocation induced by the equilibrium strategy profile corresponding to state A. So, why should players different from 1 believe the anonymous deviator when she implicitly claims that the state is 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 A. Then she would still want us to play as if we were in B, because that way she will get a better payoff. And of course, also if she does believe in B, it is better for him that we play as if we were in B. Thus she wants us to believe in B no matter what. It is as if there were no signal that 1 does not believe in A. So we will choose now what we would have chosen without any deviation from him."

The block-construction (borrowed from Wiseman 2011) aims at circumventing this kind of complications. Here are the details.

Proof

The proof consists in total of 8 steps. To give a quick overview, these are:

- Step 1: Given a target allocation we construct a sequence of allocations with utility payoffs below this allocation that will be 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 predescribed strategies is not profitable.

The details are following.

Step 1.

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

$$v_\xi^{i,n\mathrm{dev}}[\omega] := u_\xi^i \left(x_\xi^{i,n\mathrm{dev}}[\omega] \right) \ \text{ with } x_\xi^{i,n\mathrm{dev}}[\omega] := \rho_n x_\xi^{*i}[\omega] + (1 - \rho_n) w_\xi^i, \quad \rho_n \in (0,1).$$

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

$$0 < \varepsilon_n < v_\xi^{i,n\mathrm{dev}}[\omega] - v_\xi^{i,(n+1)\mathrm{dev}}[\omega] \tag{4}$$

for every player i and every state $\omega \in \Omega$. Notice that ε_n does not depend upon ξ, while the payoff $v_\xi^{i,n\mathrm{dev}}[\omega]$ does. The sequences $(\rho_n)_n$ and $(\varepsilon_n)_n$ need to be chosen so as to converge sufficiently rapidly towards 0^+ (as $n \to +\infty$) for (4) 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 trade 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 $\xi = (t, s_{t-1}, s) \in \mathbf{D}$ in period $\tau(\xi) = t \in \mathbb{N}$, a δ-action is defined as follows: Let $\delta > 0$ be small. Define the actions on the goods markets by

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

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

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

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

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

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 $(F1\xi)$ and $(F2\xi)$ are satisfied. Condition $(F3\xi)$ 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 ξ the left-hand side

of the individual budget constraint $(*_\xi^i 1)$ is equal to

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

and the right-hand side equals

$$\sum_{\ell=1}^{L} p_{\xi,\ell} q_{\xi,\ell}^i + \sum_{j=1}^{J} \pi_{\xi,j} \gamma_{\xi,j}^i + \sum_{j=1}^{J} \left(\theta_{\xi^-,j}^i - \varphi_{\xi^-,j}^i \right) D_{\xi,j}$$

$$= \sum_{\ell\in\mathcal{L}_c} \frac{N\bar{b}_\ell}{\delta} \frac{\delta}{N} + L_a 1 \frac{\delta}{N} + J1 \frac{\delta}{N}$$

$$= \sum_{\ell\in\mathcal{L}_c} \bar{b}_\ell + \frac{L_a\delta}{N} + \frac{J\delta}{N}.$$

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 $(*_\xi^i 2)$ 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 $(F2\xi)$. 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 $(*_\xi^i 1)$ 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 $(F2\xi^-)$ —where ξ^- is the predecessor of node ξ. 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 ξ, there exists a lower bound on the bids $\bar{b}_\ell$ 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 $\bar{b}_\ell$.

<u>Step 3.</u>

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 ξ, 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 $(*^i_\xi 1)$.

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

    - on the commodity markets

$$b^i_{\xi,\ell} := \begin{cases} \sum_{j=1}^J \theta^i_{\xi^-,j} C_{j\ell} & \text{for } \ell \in \mathcal{L}_c \\ \frac{N}{\delta} & \text{for } \ell \in \mathcal{L}_a \end{cases}$$

$$q^i_{\xi,\ell} := \begin{cases} \sum_{j=1}^J \varphi^i_{\xi^-,j} C_{j\ell} & \text{for } \ell \in \mathcal{L}_c \\ \frac{\delta}{N} & \text{for } \ell \in \mathcal{L}_a \end{cases}$$

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

    - on the asset markets

$$\beta^i_{\xi,j} := \frac{\delta}{N},$$
$$\gamma^i_{\xi,j} := \frac{\delta}{N}$$

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

The resulting prices are as follows:

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

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 $(F1\xi)$, $(F2\xi)$ and $(F3\xi)$. For the individual budget constraint $(*^i_\xi 1)$ at node ξ we obtain for the left-hand side

$$\sum_{\ell=1}^L b^i_{\xi,\ell} + \sum_{j=1}^J \beta^i_{\xi,j} = \sum_{\ell \in \mathcal{L}_c} \sum_{j=1}^J \theta^i_{\xi^-,j} C_{j\ell} + \frac{L_a N}{\delta} + \frac{J\delta}{N}$$

30

and for the right-hand side

$$\sum_{\ell=1}^{L} p_{\xi,\ell} q_{\xi,\ell}^{i} + \sum_{j=1}^{J} \pi_{\xi,j} \gamma_{\xi,j}^{i} + \sum_{j=1}^{J} \left(\theta_{\xi-,j}^{i} - \varphi_{\xi-,j}^{i} \right) D_{\xi,j}$$

$$= \sum_{\ell \in \mathcal{L}_c} 1 \sum_{j=1}^{J} \varphi_{\xi-,j}^{i} C_{j\ell} + L_a \frac{N^2}{\delta^2} \frac{\delta}{N} + J 1 \frac{\delta}{N} + \sum_{j=1}^{J} \left(\theta_{\xi-,j}^{i} - \varphi_{\xi-,j}^{i} \right) \left(\sum_{\ell \in \mathcal{L}_c} C_{j\ell} \right)$$

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

Moreover condition $(*_{\xi}^{i} 2)$ 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 individuals play actions (according to Lemma 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 $x_{\xi}^{*i}[\omega]$, together with actions, $\varphi_{\xi}^{*i}[\omega]$ and $\theta_{\xi}^{*i}[\omega]$, 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 ξ 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 (F2ξ). Thus, player 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 addition sale of commodities. More precisely,

- on the commodity markets, play:

$$b^i_{\xi,\ell} := \begin{cases} \frac{1}{\mathcal{L}} \sum_{j=1}^J \pi^*_{\xi,j}[\omega] \left(\varphi^{*i}_{\xi,j}[\omega] - \theta^{*i}_{\xi,j}[\omega] \right) & \text{if } \sum_{j=1}^J \pi^*_{\xi,j}[\omega] \left(\varphi^{*i}_{\xi,j}[\omega] - \theta^{*i}_{\xi,j}[\omega] \right) \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

$$q^i_{\xi,\ell} := \begin{cases} -\delta \frac{1}{\mathcal{L}} \sum_{j=1}^J \pi^*_{\xi,j}[\omega] \left(\varphi^{*i}_{\xi,j}[\omega] - \theta^{*i}_{\xi,j}[\omega] \right) & \text{if } \sum_{j=1}^J \pi^*_{\xi,j}[\omega] \left(\varphi^{*i}_{\xi,j}[\omega] - \theta^{*i}_{\xi,j}[\omega] \right) \leq 0 \\ 0 & \text{otherwise} \end{cases}$$

- on the asset markets:

$$\beta^i_{\xi,j} := \pi^*_{\xi,j}[\omega] \big(\theta^{*i}_{\xi,j}[\omega] - \eta^i_{\xi,j} \big),$$
$$\gamma^i_{\xi,j} := \varphi^{*i}_{\xi,j}[\omega] - \eta^i_{\xi,j}$$

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

Since $x^{*i}_\xi[\omega]$ is feasible, the asset markets clear, $\sum_{i=1}^N \varphi^{*i}_\xi[\omega] = \sum_{i=1}^N \theta^{*i}_\xi[\omega]$. Therefore,

$$0 = \sum_{i=1}^N \sum_{j=1}^J \pi^*_{\xi,j}[\omega] \left(\varphi^{*i}_{\xi,j}[\omega] - \theta^{*i}_{\xi,j}[\omega] \right)$$

$$= \sum_{\substack{i=1, \\ \varphi^{*i}_{\xi,j}[\omega] - \theta^{*i}_{\xi,j}[\omega] > 0}}^N \sum_{j=1}^J \pi^*_{\xi,j}[\omega] \left(\varphi^{*i}_{\xi,j}[\omega] - \theta^{*i}_{\xi,j}[\omega] \right) + \sum_{\substack{i=1, \\ \varphi^{*i}_{\xi,j}[\omega] - \theta^{*i}_{\xi,j}[\omega] < 0}}^N \sum_{j=1}^J \pi^*_{\xi,j}[\omega] \left(\varphi^{*i}_{\xi,j}[\omega] - \theta^{*i}_{\xi,j}[\omega] \right)$$

Hence,

$$\sum_{\substack{i=1, \\ \varphi^{*i}_{\xi,j}[\omega] - \theta^{*i}_{\xi,j}[\omega] > 0}}^N \sum_{j=1}^J \pi^*_{\xi,j}[\omega] \left(\varphi^{*i}_{\xi,j}[\omega] - \theta^{*i}_{\xi,j}[\omega] \right) = - \sum_{\substack{i=1, \\ \varphi^{*i}_{\xi,j}[\omega] - \theta^{*i}_{\xi,j}[\omega] < 0}}^N \sum_{j=1}^J \pi^*_{\xi,j}[\omega] \left(\varphi^{*i}_{\xi,j}[\omega] - \theta^{*i}_{\xi,j}[\omega] \right).$$

The resulting prices are as follows:

$$p_{\xi,\ell} = \frac{1}{\delta}$$
$$\pi_{\xi,j} = \pi^*_{\xi,j}[\omega].$$

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 $(F1\xi)$, $(F2\xi)$ and $(F3\xi)$. For the asset trades note that $(F1\xi)$ holds, as $x_\xi^{*i}[\omega]$ is feasible.

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

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

and for the right-hand side:

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

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

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

and for the right-hand side:

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

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 $(F2\xi)$.

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, $\mathbb{P}^i_{\xi \setminus \bar{\xi}}(h^i_\xi)$[25], 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, $\mathbb{P}^i_{\xi \setminus \bar{\xi}}(h^i_\xi)$ (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.[26] For the remaining T periods of the block, players start with a post-M-transition actions, and then play the profile that results in a stage-payoff $u^i_\xi \left(x^{*i}_\xi[\hat{\omega}] \right)$ for $\tau(\xi) = M + 2, ..., M + T$ in state $\hat{\omega}$. The actions are constructed using Lemma 2 for every node ξ with $\tau(\xi) = M + 3, ..., M + T$. 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 3 and 6.

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 (consisting in P_{n+1} blocks)

---

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

[26]Ties can be broken by some arbitrary rule.

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 a post-M-transition actions, and finally there are $T - 2$ periods where action profiles are played, such that the target allocation after the nth deviation is reached. The target allocation in the $T - 2$ last periods of a post-deviation block consists in playing a certain sequence of SSIRF 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 $v_\xi^{i,\text{1dev}}[\hat{\omega}]$ in state $\hat{\omega}$ is played, for ξ with $\tau(\xi) = M + 3, ..., M + T$. The first two periods after the M block is are post-M-transition actions. Compared to an equilibrium block, a post-deviation block consists as well 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 SSIRF allocations that are strictly worse than the target allocations of the equilibrium block or the previous post-deviation block.

Step 6.

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 nth deviation) with no unilateral deviation is followed by a post-deviation-n block.

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

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

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:

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

If M, the length of the experimentation block, is large enough, the probability that, for every i, $\mathbb{P}_{\xi \backslash \bar{\xi}}^i(h_\xi^i)$ 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 $v_\xi^{*i}[\omega]$) 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 $v_\xi^{*i}[\omega]$ and the actual payoff that accrues to player i at node ξ can be made lower than ε (for δ sufficiently small). Let us denote by $U_{\mathbf{D}(\xi_0)}^i(\sigma^*, \omega)$ the final overall payoff induced by the equilibrium strategy, and by $U_{\mathbf{D}(\xi_0)}^i(x^{*i}, \omega)$, the final payoff induced by our equilibrium target allocation.[27]

---

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

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 $\overline{u}$, while during the targeting phase (of length $T - 2$), it is ε. Suppose $T - 2 = QM$ for some integer Q. One has:

$$U^i_{\mathbf{D}(\xi_0)}(x^{*i}, \omega) - U^i_{\mathbf{D}(\xi_0)}(\sigma^*, \omega) \le (1 - \lambda) \sum_{j=0}^{+\infty} \lambda^{jQM} \Big[\frac{1 - \lambda^{M+2}}{1 - \lambda} \overline{u} + \lambda^{M+2} \frac{1 - \lambda^{QM}}{1 - \lambda} \varepsilon \Big],$$

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,

$$U^i_{\mathbf{D}(\xi_0)}(x^{*i}, \omega) - U^i_{\mathbf{D}(\xi_0)}(\sigma^*, \omega) \le \frac{1 - \lambda^{M+2}}{1 - \lambda^{QM}} \overline{u} + \lambda^{M+2} \varepsilon.$$

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

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.[28] By manipulating allocations (hence stage-payoffs), a player cannot prevent her opponents from eventually learning the true state of the world, ω.

<u>Step 8.</u>

---

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

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 $\xi' = (t', s_{t'-1}, s')$, that it is the $(n+1)$th deviation observed during the play and that there are no further deviations at later nodes. It will at most yield $\overline{u}^i$ to player i. Then, the post-deviation payoff can be made ε-close to the following maximum:

$$
\begin{aligned}
(1-\lambda)\lambda^{t'-1}\Big[\overline{u}^i &+ \sum_{t=2}^{M+T} \lambda^{t-1} E_{\mathbf{P}^i_{\xi'}(\sigma)}[u^i_\xi\left(w^i_\xi\right)] \\
&+ \sum_{k=1}^{P_{n+1}} \Big[\sum_{t=k(M+T)+1}^{(k+1)(M+T)} \lambda^{t-1} E_{\mathbf{P}^i_{\xi'}(\sigma)}[u^i_\xi\left(w^i_\xi\right)] \\
&+ \sum_{k \geq P_{n+1}+1} \Big[\sum_{t=k(M+T)+1}^{k(M+T)+M} \lambda^{t-1} E_{\mathbf{P}^i_{\xi'}(\sigma)}[u^i_\xi\left(w^i_\xi\right)] \\
&+ \sum_{t=k(M+T)+M+1}^{k(M+T)+M+2} \lambda^{t-1}\overline{u} \\
&+ \sum_{t=k(M+T)+M+3}^{(k+1)(M+T)} \lambda^{t-1} E_{\mathbf{P}^i_{\xi'}(\sigma)}[v^{i,(n+1)\mathrm{dev}}_\xi[\omega]]\Big] \Big].
\end{aligned} \tag{5}
$$

Indeed, the long-run discounted payoff after a deviation consists of once 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, $E_{\mathbf{P}^i_\xi(\sigma)}[v^{i,(n+1)\mathrm{dev}}_\xi[\hat{\omega}]]$, computed with the most likely state, $\hat{\omega}$ (according to the players' truncated belief), can also be made arbitrarily close to $E_{\mathbf{P}^i_\xi(\sigma)}[v^{i,(n+1)\mathrm{dev}}_\xi[\omega]]$.

By contrast, if the $(n+1)$th deviation did <u>not</u> 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:

$$(1-\lambda)\lambda^{t'-1}\sum_{k\geq 0}\Big[\sum_{t=k(M+T)+1}^{k(M+T)+M}\lambda^{t-1}E_{\mathbf{P}^i_{\xi'}(\sigma)}[u^i_\xi\left(w^i_\xi\right)]$$

$$+\sum_{t=k(M+T)+M+3}^{(k+1)(M+T)}\lambda^{t-1}E_{\mathbf{P}^i_{\xi'}(\sigma)}[v^{i,n\mathrm{dev}}_\xi[\omega]]\Big]. \tag{6}$$

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

In order to check whether the difference (6) - (5) is positive, all we need is to ensure that:

$$(1-\lambda)\lambda^{t'-1}\Big[\underline{u}-\overline{u}+\sum_{k=1}^{P_{n+1}}\sum_{t=k(M+T)+M+3}^{(k+1)(M+T)}\lambda^{t-1}E_{\mathbf{P}^i_{\xi'}(\sigma)}\Big[v^{i,n\mathrm{dev}}_\xi[\omega]-u^i_\xi\left(w^i_\xi\right)\Big]$$

$$+\varepsilon_n\sum_{k\geq P_{n+1}+1}\Big[\sum_{t=k(M+T)+M+3}^{(k+1)(M+T)}\lambda^{t-1}\Big]-\overline{u}\sum_{k\geq 0}\Big[\sum_{t=k(M+T)+M+1}^{k(M+T)+M+2}\lambda^{t-1}\Big]\Big]>0$$

Note that since $v^{i,n\mathrm{dev}}[\omega]$ results from a SSIRF allocation, we have

$$E_{\mathbf{P}^i_{\xi'}(\sigma)}[v^{i,n\mathrm{dev}}_\xi[\omega]-u^i_\xi\left(w^i_\xi\right)]>0$$

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

$$g_\xi:=\min_{i\in\mathcal{N}}E_{\mathbf{P}^i_{\xi'}(\sigma)}[v^{i,n\mathrm{dev}}_\xi[\omega]-u^i_\xi\left(w^i_\xi\right)].$$

It is sufficient to require that:

$$(1-\lambda)\lambda^{t'-1}\Big[\underline{u}-\overline{u}+\sum_{k=1}^{P_{n+1}}\Big[\sum_{t=k(M+T)+M+3}^{(k+1)(M+T)}\lambda^{t-1}g_\xi\Big]$$

$$+\varepsilon_n\frac{(1-\lambda^{T-2})}{(1-\lambda)(1-\lambda^{M+T})}\lambda^{(P_{n+1}+1)(M+T)+M+2}$$

$$-\overline{u}\frac{(1-\lambda^2)}{(1-\lambda)(1-\lambda^{M+T})}\lambda^M\Big]>0,$$

It is easy to see that, whatever being the distance, $(1-\lambda)\underline{u}-\left((1-\lambda)+\frac{(1-\lambda^2)}{(1-\lambda^{M+T})}\lambda^M\right)\overline{u}$, and for every $\varepsilon_n>0$, and every $\lambda\geq\lambda_\varepsilon$, there exists some integer $P^{\lambda,\varepsilon}_{n+1}$ 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 (5).

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 $\forall n$, $P_n \geq P_n^{\lambda,\varepsilon}$.

$\square$

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 lead to additional leverage and asset price volatility (relative to a model with homogeneous beliefs or relative to equilibria in the complete markets economy). Our result suggests that this conclusion is partly due to his narrow (though standard) definition of perfect competition. Indeed, due to imperfect competition, those traders with incorrect beliefs can strategically *learn* the state of the world. We therefore 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.

Let us end with a final remark concerning the link of the present work with the perfectly competitive set-up. In Giraud and Weyers (2004), as already mentioned, a first step towards the present Folk theorem had been obtained in the particular setting of *exogenously* incomplete markets (with finite horizon). Here, we get a Folk theorem for economies where missing markets are *endogenously* determined, due both to the presence of collateral constraints and to the lack of complete information. At the end of Giraud and Weyers (2004), however, the asymptotic properties of type-symmetric strategic equi-

libria were studied when the number of individuals of each type grows to infinity. It was shown that there is a discontinuity at the limit: Indeed, the limit-set of equilibria remains quite large while it is well-known that, at least with real assets, finite-horizon economies with incomplete markets generically admit a finite number of perfectly competitive equilibria (Duffie and Shafer 1985). An analogous remark holds in the present incomplete information set-up. Suppose that each type of player is actually represented by K identical individuals, and let $K \to +\infty$. The same argument as in Giraud and Weyers (2004) allows us to extend our partial Folk theorem to the asymptotic case. Therefore, we get that, at the limit, there is still a continuum of Bayesian perfect equilibria, exhibiting a large variety of efficiency properties (although each individual is negligible). It also suggests that, despite the considerable literature devoted to its foundation, the very concept of perfect competition itself deserves further investigation. In particular, whether it is captured as a price-taking assumption or else as the limit benchmark obtained by letting the weight of each price-maker shrink to zero does not lead to the same conclusion.

Our result and this last observation suggest that considerable care is necessary in invoking the impact of collateral regulation on the inefficiency of equilibria with private information —both in perfectly and imperfectly competitive environments.

References

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.

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500.

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.

Araujo, A., Pascoa, M. R., and Torres-Martinez, J. P. (2002). Collateral avoids Ponzi schemes in incomplete markets. *Econometrica*, 70(4):1613–1638.

Arnott, R. and Stiglitz, J. E. (1986). Moral hazard and optimal commodity taxation. *Journal of Public Economics*, 29(1):1–24.

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.

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.

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.

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.

Brangewitz, S. (2011). Finite horizon strategic market games with collateral. Unpublished.

Cao, D. (2011). Collateral shortages, asset price and investment volatility with heterogeneous beliefs. Unpublished.

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.

Dubey, P. and Geanakoplos, J. (2003). Monetary equilibrium with missing markets. *Journal of Mathematical Economics*, 39(5-6):585–618.

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.

Facchinei, F. and Kanzow, C. (2010). Generalized Nash equilibrium problems. *Annals of Operations Research*, 175(1):177–211.

Florenzano, M. and Gourdel, P. (1996). Incomplete markets in infinite horizon: Debt constraints versus node prices. *Mathematical Finance*, 6(2):167–196.

Friedman, M. (1953). *Essays in Positive Economics*. University of Chicago Press.

Geanakoplos, J. and Zame, W. R. (2007). Collateralized asset markets. *UCLA Working Paper*.

Giraud, G. (2003). Strategic market games: an introduction. *Journal of Mathematical Economics*, 39:355–375.

Giraud, G. and Weyers, S. (2004). Strategic market games with a finite horizon and incomplete markets. *Economic Theory*, 23:467–491.

González-Díaz, J. and Meléndez-Jiménez, M. A. (2011). On the notion of perfect Bayesian equilibrium. Unpublished.

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.

Harker, P. T. (1991). Generalized Nash games and quasi-variational inequalities. *European Journal of Operational Research*, 54(1):81–94.

Heifetz, A. (2006). The positive foundation of the common prior assumption. *Games and Economic Behavior*, 56(1):105–120.

Magill, M. and Quinzii, M. (1994). Infinite horizon incomplete markets. *Econometrica*, 62(4):853–880.

O'Connell, S. A. and Zeldes, S. P. (1988). Rational Ponzi games. *International Economic Review*, 29(3):431–450.

Peck, J., Shell, K., and Spear, S. E. (1992). The market game: Existence and structure of equilibrium. *Journal of Mathematical Economics*, 21:271–299.

Postlewaite, A. and Schmeidler, D. (1978). Approximate efficiency of non-Walrasian Nash equilibria. *Econometrica*, 46(1):127–135.

Prescott, E. C. and Townsend, R. M. (1984a). General competitive analysis in an economy with private information. *International Economic Review*, 25(1):1–20.

Prescott, E. C. and Townsend, R. M. (1984b). Pareto optima and competitive equilibria with adverse selection and moral hazard. *Econometrica*, 52(1):21–46.

Radner, R. (1979). Rational expectations equilibrium: Generic existence and the information revealed by prices. *Econometrica*, 47(3):655–678.

Rahi, R. (1995). Optimal incomplete markets with asymmetric information. *Journal of Economic Theory*, 65(1):171–197.

Rothschild, M. (1974). A two-armed bandit theory of market pricing. *Journal of Economic Theory*, 9(2):185–202.

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.

Sandroni, A. (2000). Do markets favor agents able to make accurate predictions? *Econometrica*, 68(6):1303–1341.

Shapley, L. S. and Shubik, M. (1977). Trade using one commodity as a means of payment. *Journal of Political Economy*, 85:937–968.

Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–374.

Thomas, C. P. (1995). The role of fiscal policy in an incomplete markets framework. *The Review of Economic Studies*, 62(3):449–468.

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.

Wiseman, T. (2005). A partial folk theorem for games with unknown payoff distributions. *Econometrica*, 73(2):629–645.

Wiseman, T. (2011). A partial folk theorem for games with private learning. Forthcoming in Theoretical Economics.

Zhu, H. (2011). Do dark pools harm price discovery? Available at SSRN: http://ssrn.com/abstract=1712173.

A Appendix

A.1 Proof of Lemma 1

To show Lemma 1 for our model, we modify the proof of Giraud and Weyers (2004) slightly.

Proof. Fix a node $\xi \in \mathbf{D}$ at time $t \leq T$. Since the allocation of initial endowments $(w_\xi^i)_i$ are Pareto-inefficient in the L-good spot economy there exists a consumption stream $(\bar{x}_\xi^i)_i$ that Pareto dominates $(w_\xi^i)_i$ and satisfies for every good $\ell \in \mathcal{L}$

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

By the strict monotonicity of the preferences, there exists a consumption stream $(\bar{x}'^i_\xi)_i$ such that

$$u_\xi^i(\bar{x}'^i_\xi) > u_\xi^i(\bar{x}_\xi^i) \quad i = 1, ..., N$$

and

$$\sum_{i=1}^N \bar{x}'^i_{\xi,\ell} = \sum_{i=1}^N w_{\xi,\ell}^i.$$

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

$$p_\xi \cdot \bar{x}'^i_\xi = p_\xi \cdot w_\xi^i$$

is satisfied and furthermore

$$E_\omega[u_\xi^i(\bar{x}'^i_\xi)] > E_\omega[u_\xi^i(\bar{x}_\xi^i)] \geq E_\omega[u_\xi^i(w_\xi^i)]$$

for all $i \in \mathcal{N}$ and $t \leq T$. □

A.2 Proof of Lemma 2

To show Lemma 2 for our model, we modify the proof of Giraud and Weyers (2004).

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

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

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

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

$$\sum_{i=1}^{N} q^i_{\xi,\ell} = \sum_{i=1}^{N} \left(w^i_{\xi,\ell} + \sum_{j=1}^{J} \bar{\varphi}^i_{\xi^-,j} C_{j\ell} \right)$$

$$= \sum_{i=1}^{N} \left(\bar{x}^i_{\xi,\ell} + \sum_{j=1}^{J} \bar{\varphi}^i_{\xi,j} C_{j\ell} \right),$$

$$\sum_{i=1}^{N} b^i_{\xi,\ell} = p_{\xi,\ell} \sum_{i=1}^{N} \left(\bar{x}^i_{\xi,\ell} + \sum_{j=1}^{J} \bar{\varphi}^i_{\xi,j} C_{j\ell} \right).$$

Hence,

$$\bar{p}_{\xi,\ell} = \frac{\sum_{i=1}^{n} b^i_{\xi,\ell}}{\sum_{i=1}^{n} q^i_{\xi,\ell}}$$

$$= p_{\xi,\ell}.$$

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

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

$$\pi_{\xi,j} = \frac{\sum_{i=1}^{N} \beta^i_{\xi,j}}{\sum_{i=1}^{N} \gamma^i_{\xi,j}}$$

$$= \frac{\bar{\pi}_{\xi,j} \sum_{i=1}^{N} \bar{\theta}^i_{\xi,j}}{\sum_{i=1}^{N} \bar{\varphi}^i_{\xi,j}}$$

$$= \bar{\pi}_{\xi,j}$$

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

$$\pi_{\xi,j} = \frac{\sum_{i=1}^{N} \beta_{\xi,j}^i}{\sum_{i=1}^{N} \gamma_{\xi,j}^i}$$

$$= \frac{\bar{\pi}_{\xi,j} \sum_{i=1}^{N} \frac{\delta}{N}}{\sum_{i=1}^{N} \frac{\delta}{N}}$$

$$= \bar{\pi}_{\xi,j}.$$

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

$$\varphi_{\xi,j}^i = \gamma_{\xi,j}^i,$$

$$\theta_{\xi,j}^i = \frac{\beta_{\xi,j}^i}{\pi_{\xi,j}}.$$

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

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

Therefore,

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

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

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

It remains to check that the budget constraint $(\ast_{\xi}^i 1)$ for the bids and offers is satisfied.

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

Inserting the assumed aci for $b_{\xi,\ell}^i$, $q_{\xi,\ell}^i$, $\gamma_{\xi,j}^i$ and $\beta_{\xi,j}^i$ we obtain for $(\ast_{\xi}^i 1)$

47

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

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

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

which holds since $(\bar{x}^i, \bar{\varphi}^i, \bar{\theta}^i)_{i \in \mathcal{N}}$ was assumed to be a feasible allocation.

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

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

As $(w_\xi^i)_i \gg 0$, this strategy profile is full. This completes the proof. $\qquad\square$