REPEATED GAMES WITH INCOMPLETE INFORMATION AND DISCOUNTING

MARCIN PĘSKI

ABSTRACT. We analyze discounted repeated games with incomplete information, such that the players' payoffs depend only on their own type (*known-own payoff case*). We describe an algorithm for finding all equilibrium payoffs in games for which there exists an open set of belief-free equilibria of Hörner and Lovo [2009]. This includes generic games with one-sided incomplete information and a large and important class of games with multi-sided incomplete information. When players become sufficiently patient, all Bayesian Nash equilibrium payoffs can be approximated by payoffs in sequential equilibria in which information is revealed finitely many times. The set of equilibrium payoffs is typically larger than the set of equilibrium payoffs in repeated games without discounting, and larger than the set of payoffs obtained in belief-free equilibria. The results are illustrated in bargaining and oligopoly examples.

1. INTRODUCTION

This paper contributes to the literature on repeated games with discounting and incomplete information, in which players know their own payoffs. It introduces a payoff set based on two geometric operations, and establishes two results. First, all elements of this set are (sequential) equilibrium payoff vectors when players are sufficiently patient. Second, for a rich class of games (those in which belief-free equilibria exist, in particular, games with one-sided incomplete information), it proves that this characterization is tight: no payoff vector outside this set can be achieved in a (Bayes) Nash equilibrium.

Date: 03/08/13.

I am grateful to Françoise Forges, Tom Wiseman, Johannes Hörner, and anonymous referees for helpful comments. All the remaining errors are my own.

One of the major problems in the analysis of repeated games with incomplete information is that there is no natural candidate for the payoff set. This problem is not present in games with complete information, where it is immediately shown that all equilibrium payoffs must be feasible and individually rational, and the main difficulty is to find conditions under which all feasible and individually rational payoffs can be attained in subgame perfect equilibria. With incomplete information, the set of (naturally defined) feasible and individually rational payoffs is typically too large, since not all such payoffs can be attained (or even approximated) by equilibrium payoffs. Another candidate, the equilibrium payoff set obtained in the literature without discounting (Hart [1985], Shalev [1994], and Koren [1992]), is typically smaller than the set of payoffs that can be obtained in games with discounting.

We solve the problem in a two-part argument. First, we construct a candidate equilibrium correspondence that assigns a payoff set to each prior belief. The idea is to consider payoffs in strategy profiles in which (i) there are finitely many periods in which players reveal information (by taking partially or fully separating actions), (ii) these periods are separated by possibly long time intervals during which the types of players pool their actions, and (iii) at each period, the continuation payoffs are individually rational. The construction begins with a set of individually rational payoffs in profiles in which no information is ever revealed. Next, we go through a sequence of steps that alternate between two geometric operations: (A) for each prior, the payoff set is convexified with the payoffs in profiles in which the players do not reveal any information, and (B) the payoffs are convexified across different initial priors. Operation (A) constructs profiles in which during the initial periods, the types of players pool their actions, and in operation (B), we obtain profiles in which the types reveal some substantial information in the first period. Operations (A) and (B) correspond to two properties of bi-martingales from the literature on games without discounting with a key difference: because the initial play does not matter in that literature, the analogue of operation (A) replaces the set of payoffs with its convex hull, but it does not convexify it with the non-revealing payoffs. For sufficiently patient players, all the payoff vectors in the candidate equilibrium correspondence can be attained by payoffs in *finitely revealing equilibria*: sequential equilibria in which players' information is revealed at most finitely many times.

Second, we show that no payoff outside the constructed set can be attained in equilibrium. For the second part, we assume that there exists an open set of payoffs in *belief-free equilibria* of Hörner and Lovo [2009]: sequential equilibria in which during the first period of the game, all players fully reveal their information (i.e., they take fully separating actions), and such that the players are *ex post* indifferent between revealing their type truthfully or reporting any other type (i.e., they are indifferent conditional on any type of the opponent). The payoffs in such equilibria form a multilinear cross section (i.e., a *thread*) of payoffs across games with all possible initial beliefs. The assumption is generically satisfied in games with one-sided incomplete information (like oligopoly models). However, there are robust examples of games with two players and incomplete information on both sides that do not satisfy the assumption.

Given this assumption, we show that *all* payoffs attained in Nash equilibria of the repeated game can be approximated by the payoffs constructed in the first part of our argument, i.e., by payoffs in finitely revealing equilibria. The idea is to modify the equilibrium continuation payoffs to pull them toward the thread of payoffs in the belief-free equilibria. We show that the continuation payoffs reach the neighborhood of the belief-free equilibrium payoff set in finitely many periods. Once this happens, we finish the construction with an immediate and full revelation of information. The argument is of independent interest, as it is very simple and possibly can be applied in other related settings (like games with types that may slowly change over time).

At first sight, the characterization of all equilibrium payoffs through finitely revealing profiles has a simple intuition. Because the beliefs of the players are martingales, they converge. This means that with a high probability, players stop revealing substantial amounts of information after finitely many periods. Nevertheless, the result is far from obvious. The above intuition does not lead to the proof as it leaves open the possibility that after low probability histories, the continuation game requires a large amount of information to be revealed. In fact, the intuition fails in repeated

games with no discounting in which there are examples of games with equilibrium payoffs that cannot be approximated by finitely revealing profiles (see Forges [1984], Forges [1990] and Aumann and Hart [2003]).

The previous literature has not been able to describe the equilibrium set except for very special cases. The closest to the current paper is Peski [2008], who characterized the equilibrium payoffs in games with incomplete information only on one side, and with the informed player having only two types. The current paper generalizes Peski [2008] to multiple types and multi-sided incomplete information. The characterization of payoffs in finitely revealing equilibria is a relatively straightforward extension of Peski [2008]. The key step of the current paper, i.e., the argument that no other payoffs can be attained in equilibrium, is entirely novel.¹

The main advantage of our characterization is that the payoff set in finitely revealing profiles can be computed. We illustrate this claim with two examples. First, we discuss a class of oligopoly games. That class includes, as a knife-edge case, a Bertrand oligopoly with privately known production costs from Athey and Bagwell [2008]. In that paper, the authors propose mechanism design methods for analyzing repeated games with incomplete information. They describe the equilibrium that maximizes the sum of *ex ante* payoffs among all symmetric equilibria, and they show that no information is revealed in such an equilibrium. Here, we explain that there is a relation between the mechanism design approach and equilibria in which all players fully and immediately reveal their information. We show that in oligopoly games, all equilibrium payoffs can be approximated by such equilibria, and can be derived as solutions to a simple mechanism design problem. We use the explicit description of payoffs to show that some (and, in some cases, complete) productive efficiency can typically be obtained in the Pareto-dominant equilibrium. In particular, we argue

¹To compare, notice that Peski [2008] uses a much more complicated differential technique that, despite our best efforts, could not be extended beyond the two-type, one-sided case. In addition, notice that Athey and Bagwell [2008] (see later in the introduction) solve the second part of the argument (i.e., showing that there are no other Nash equilibria) using a sophisticated approach from the mechanism design literature and by making assumptions about the log-concavity of the cost distribution. Their approach does not seem to generalize well beyond the particular example they analyze.

that the "pooling" result of Athey and Bagwell [2008] is not robust to alternative demand specifications.

Second, we discuss a bargaining game with two players, one-sided incomplete information, and two types (normal and "strong") of the informed player. We assume that the game between the normal type and the uninformed player has strictly conflicting interests (Schmidt [1993]). The strong type's payoffs are parametrized as a convex combination between the payoffs of the normal type and the payoffs of a player who is committed to playing a single action (i.e., for whom repeating a single action is a dominant strategy in the repeated game). We describe the Pareto frontier of the equilibrium set as a solution to a system of differential equations. We show that there are efficient equilibria that require any arbitrarily large number of periods with the information revelation. When the payoffs of the strong type converge to the payoffs of the committed player, all equilibrium payoffs converge to the Stackelberg outcome of the informed player.

We compare our characterization to the literature on repeated games without discounting (see Aumann and Maschler [1995]). That literature typically considers the general payoff case, in which players' payoffs may depend on the types of their opponents.² Hart [1985] considers one-sided uncertainty with general payoffs and characterizes the equilibrium payoffs as values of bi-martingales. The characterization in the general-payoff no-discounting case is not constructive and there is no known algorithm for finding all equilibrium payoffs in the general case. This fact is related to the existence of games and equilibria that cannot be approximated by payoffs in the finitely revealing equilibrium payoffs can be obtained with strategies in which players reveal all their information in the first period (Shalev [1994] and Koren [1992]). Some of the similarities and the differences between the no-discounting and discounted cases are discussed in more detail in section 7.

 $^{^{2}}$ To avoid players learning about the other players' types from their own payoffs, the literature assumes that the payoffs are not observed until the end of the (infinite) repeated game. This assumption is not needed in the known-own payoffs case.

An important difference between the discounted and no-discounting case is that equilibria always exist in the former case, but not necessarily in the latter. However, the assumption made in this paper (i.e., the set of belief-free payoffs has a non-empty interior) ensures that the set of payoffs in the repeated game without discounting has a non-empty interior. The same assumption ensures that our candidate set of payoffs is non-empty (as it contains the set of equilibrium payoffs in games without discounting). In general, the non-emptiness of the candidate payoff sets, as well as the characterization of the equilibrium payoffs without the assumption, remains an open question.

Kreps and Wilson [1982a] and Milgrom and Roberts [1982] introduced a model of reputation with one-sided incomplete information about the type of the long-run informed player: strategic or commitment ("reputational") types. This literature was extended to equal discounting and patient players in Cripps and Thomas [1997], Chan [2000], and Cripps et al. [2005]. Because in the reputational model, the highest payoff of the commitment type is equal to his minmax payoffs, this model does not have an open set of payoffs. However, a small perturbation of the reputational types' payoffs may create an open thread and restore the assumption. We can use "nearby" models to test the predictions of the reputational literature. Our third example illustrates the robustness of the result of Cripps et al. [2005].

One of the first papers to study repeated games with one-sided incomplete information and equal discounting is Cripps and Thomas [2003] (see also Bergin [1989]). They look at the limit correspondence of payoffs when the probability of one of the types is close to 1.³ They show that the set of payoffs of the uninformed player and the high probability type is close to the folk theorem payoff set from a complete information game. Cripps and Thomas [1997] and Chan [2000] ask similar questions within the framework of reputation games. All these results are proven by the construction of finitely revealing equilibria.

³Cripps and Thomas [2003] also discuss the limit of payoff sets when the two players become infinitely patient, but player I becomes patient much more quickly than player U. Their characterization is closely related to Shalev and Koren's results for the no-discounting case.

Hörner and Lovo [2009] study the general payoff case with multi-sided incomplete information and characterize the set of payoffs obtained in *belief-free* equilibria. Hörner et al. [2011] describe detailed conditions for information structures in N-player games under which belief-free equilibria exist for *all* payoff functions. Our main result is limited to games in which belief-free equilibria exist. However, we characterize the set of *all* equilibrium payoffs and not just the set of payoffs in belief-free equilibria. In particular, even if belief-free equilibria exist, they may not capture all equilibrium payoffs, or not even all efficient equilibrium payoffs. (See examples at the end of section 6.1 and in section 6.2.)

There are other related papers on repeated games with discounting but with different kinds of incomplete information. Wiseman [2005] considers the situation in which the payoffs are initially unknown by all players (i.e., there is no asymmetric incomplete information), and the players learn the payoff function from observing the realization of their payoffs over time. Fudenberg and Yamamoto [2010] and Fudenberg and Yamamoto [2011] study the case where the payoffs and the monitoring structure are initially unknown, and the players may start the game with private information about the state of the world. The players learn over time by observing signals. The authors find conditions on the informativeness of the signals that ensure that in equilibrium, players can learn the state very quickly, and the set of equilibrium payoffs obtained in each state is equal to the folk theorem set in the complete information game *as if* the players knew the state from the beginning. In this setting, the set of payoffs is not affected by initially incomplete information.

In the next section, we describe the model and preliminary results. In section 3, we describe the geometric construction of the candidate payoff set. In section 4, we show that each element of the payoff set can be attained in finitely revealing equilibria. In section 5, we show that given the existence of an open set of payoffs in belief-free equilibria, each Nash equilibrium payoff can be approximated by a payoff in a finitely revealing equilibrium. We illustrate the result with examples in section 6. In section 7, we discuss the relation to the no-discounting literature. Section 8 concludes. Most of the proofs are postponed untill the appendices.

2. Model

2.1. Repeated game. For each set $X \subseteq \mathbb{R}^d$, we write int X, clX, and conX to denote the interior, closure, and convexification of X, respectively. For each $u \in \mathbb{R}^d$, each $\varepsilon > 0$, let $B(u, \varepsilon) = \{u' : \sup_i |u_i - u'_i| < \varepsilon\}$ be an open ball in the "sup" metric.

There are *I* players, i = 1, 2, ..., I. In each period $t \ge 0$, each player *i* takes an action a_i from the finite set A_i^4 and receives payoffs $g_i(a_i, a_{-i}, \theta_i)$. The payoffs depend on the actions of all players and on the privately known type θ_i of player *i* (known-payoff case). We assume that $|A_i| \ge |\Theta_i|$ for each player *i*. Let $M = \max_{i,a,\theta_i} |g_i(a, \theta_i)| < \infty$ be the upper bound on the absolute value of the payoffs. All actions are observed.

The type of player *i* is chosen by Nature once and for all from the finite set Θ_i and revealed to player *i* before the first period of the repeated game. We write $\Theta_{-i} = \times_{j \neq i} \Theta_j$ to denote the set of type tuples of all players but *i*, and $\Theta = \times_i \Theta_i$ to denote the set of type profiles. We also write $\Theta^* = \Theta_1 \cup ... \cup \Theta_I$ to denote the disjoint union of the sets of types for each player.

We encode the payoffs of different types of different players as a tuple $v = (v_i(\theta_i))_{i,\theta_i} \in \mathbb{R}^{\Theta^*}$, with the interpretation that $v_i(\theta_i)$ is the (expected) payoff of type θ_i of player *i*. We write $v_i \in \mathbb{R}^{\Theta_i}$ to refer to the component of *v* that consists of payoffs of player *i*'s types.⁵

Each type θ_i of player *i* starts the game with beliefs $\pi^{\theta_i} \in \Delta \Theta_{-i}$ about the distribution of the types of the other players. The beliefs may differ across types, and we do not assume that the players' beliefs are derived from a common prior. However, we

⁴The results of the paper extend to the case of infinite action sets with some modifications of the definitions. First, to avoid problems with updating on non-atomic, positive probability events, we assume that the players are restricted to mixed strategies with countable supports. Second, the (σ, π) -consistent beliefs in the sequential equilibrium must be obtained as the limits of the beliefs in the convergent *nets of* strategies $(\sigma_{\xi})_{\xi \in E}$ that converge to σ . The details are available upon request.

⁵The convention of encoding payoffs given one's own type follows Peski [2008] (see also the statement of the main result in Hart [1985]) and differs from some other papers in the literature. For example, Hörner and Lovo [2009] write $v \in R^{I \times \Theta}$ to denote the vector of the players' payoffs given the realization of the entire type profile, and not only the player's own type. Our convention is simpler and more natural in the known-own payoff case.

assume that all types of all players agree on which types have positive or zero probability. Precisely, from now on, we assume that each belief system $\pi = (\pi^{\theta_i})_{i,\theta_i \in \Theta_i}$ satisfies the *common rectangular support* property: for each player j, there exists set $\Theta_j^{\pi} \subseteq \Theta_j$ such that for each type θ_i of each player i, $\pi^{\theta_i}(\theta_{-i}) > 0$ if and only if $\theta_{-i} \in \times_{j \neq i} \Theta_j^{\pi}$. We refer to Θ_j^{π} as the π -support of player j. We say that type θ_j has π -positive probability if $\theta_j \in \Theta_j^{\pi}$, and π -zero probability otherwise. Let Π denote the space of belief systems with common rectangular support.⁶

Players discount the future with the common discount factor $\delta < 1$. We refer to the game with discount factor δ and initial beliefs π as $\Gamma(\pi, \delta)$.

For simplicity, we assume that players have access to public randomization. As is standard practice in the literature, we omit the reference to public randomization in the formal definition of a history.

Let $H_t = A^t$ be the set of *t*-period histories $h_t = (a_s)_{s=0}^{t-1}$. A (repeated game) strategy of player *i* is a mapping $\sigma_i : \Theta_i \times \bigcup_t H_t \to \Delta A_i$. For any profile $\sigma = (\sigma_i)_i$ of such strategies, let

$$v^{\pi,\delta}\left(\sigma\right) = (1-\delta) \sum_{\theta_{-i}\in\Theta_{-i}} \pi^{\theta_{i}}\left(\theta_{-i}\right) E_{\sigma\left(\theta_{i},\theta_{-i}\right)} \sum_{t} \delta^{t} g\left(a_{t},\theta_{i}\right) \in R^{\Theta^{*}}$$

denote the (normalized) expected payoff of player *i* type θ_i , where the expectation is computed with respect to the distribution over histories induced by strategies σ and given types (θ_i, θ_{-i}) . Let $v^{\pi,\delta}(\sigma) \in R^{\Theta^*}$ denote the (normalized) expected payoff vector.

2.2. Feasible, non-revealing payoffs. Two sets play an important role in our characterization. The first set consists of stage-game payoffs obtained when all types of each player pool their actions. For each action profile $a = (a_i) \in A \equiv \times_i A_i$, let $g(a) = (g_i(a, \theta_i))_{i,\theta_i} \in R^{\Theta^*}$ be the payoff vector obtained when each type of player *i* plays the same action a_i . Let

$$V = \operatorname{con} \left\{ g\left(a\right) : a \in A \right\} \subseteq R^{\Theta^*}$$

⁶Note that if the prior beliefs have the common rectangular support, then the posterior beliefs after positive probability events have it as well. Thus, the rectangularity will be preserved by consistent beliefs (see the definition below).

be the convex hull of payoff vectors g(a). We refer to V as *feasible*, *non-revealing* (i.e., pooling) *payoffs*.

The set V is typically smaller than the set of all feasible payoffs in game $\Gamma(\pi, \delta)$. The latter can be defined as the convex hull of payoff vectors $v^{\pi,\delta}(\sigma)$ for all strategy profiles σ , including profiles in which players' types do not pool their actions.

2.3. Individual rationality. The second set consists of individually rational payoffs. We follow Blackwell [1956], who solved the problem of extending individual rationality to the incomplete information case. Define the weighted minmax of player i: for each $\phi \in R^{\Theta_i}_+$

$$m_{i}(\phi) := \inf_{\alpha_{-i} \in \times_{j \neq i} \Delta A_{j}} \sup_{\alpha_{i} \in \Delta A_{i}} \sum_{\theta_{i}} \phi_{\theta_{i}} g\left(\alpha_{i}, \alpha_{-i}, \theta_{i}\right).$$
(2.1)

Define the set of *individually rational* payoffs as

$$IR = \left\{ v \in R^{\Theta^*} : \forall i \forall \phi \in R_+^{\Theta_i}, \ \phi \cdot v_i \ge m_i(\phi) \right\}.$$

Blackwell [1956] shows that for each payoff vector v from the interior of the set IR, for each player i, there exists a sufficiently large horizon T and a T-period strategy of players -i that ensures that the T-period average payoff of player i type θ_i is smaller than $v_i(\theta_i)$. For the precise statement of the result in the case of games with discounting, see Lemma 4 in Appendix B.

2.4. Equilibrium. A strategy profile σ is a (Bayesian) Nash equilibrium in game $\Gamma(\pi, \delta)$ for some $\pi \in \Pi$ if for each player *i* type θ_i , strategy $\sigma_i(\theta_i)$ is the best response of type θ_i . One shows that any payoff vector in a Nash equilibrium must belong to set IR.

A strategy profile σ is totally mixed if for each player *i*, type θ_i , history h_t , action $a_i, \sigma_i(a_i|h_t, \theta_i) > 0$. Each totally mixed strategy profile σ , together with the initial belief system $\pi \in \Pi$, induces a well-defined belief mapping $p^{(\sigma,\pi)} : \bigcup_t H_t \to \Pi$ through the Bayes formula. (Note that the posterior beliefs always have common rectangular support: see Footnote 6.) For any strategy profile σ , say that belief mapping p: $\bigcup_t H_t \to \Pi$ is $(\sigma, \pi) - consistent$, if there exists a sequence of totally mixed strategy profiles $\sigma_n \to \sigma$ such that $p^{(\sigma_n,\pi)} \to p$.⁷ If history h_t has a positive probability, i.e., if

⁷In both cases, we use the "pointwise" notion of convergence, i.e., $\sigma_n \to \sigma$ if and only if $\sigma_n(\theta_i, h) \to \sigma(\theta_i, h)$ for each type θ_i and history h. Our analysis would not be affected if instead we used the

for each player i,

$$\prod_{s < t} \sigma_i \left(a_s^i | h_s, \theta_i \right) > 0,$$

then $p(h_t)$ does not depend on the choice of sequence σ_n . We use this observation without any further reminder.

A strategy profile σ is a sequential equilibrium in game $\Gamma(\pi, \delta)$ if there exists (σ, π) consistent belief mapping p such that for each player i type θ_i , history h_t , continuation
strategy $\sigma_i(h_t, .)$ is the best response to continuation strategy $\sigma_{-i}(h_t, .)$, given beliefs $p_i(h_t)$. A sequential equilibrium is *n*-revealing if for any positive probability history h, there exists at most n periods t such that $p(h_t) \neq p(h_{t-1})$. A finitely revealing
equilibrium is a sequential equilibrium profile σ that is *n*-revealing for some n.

Let $NE^{\delta}(\pi)$, $FR_n^{\delta}(\pi) \subseteq R^{\Theta^*}$ be the sets of expected payoff vectors $v^{\pi,\delta}(\sigma)$ in Nash, and *n*-revealing equilibria σ , respectively. The sets of equilibrium payoffs typically depend on initial beliefs. Because any equilibrium play in which information is revealed has continuation play in a game with posterior beliefs that may differ from the prior, the payoff sets for different beliefs are related to each other. Thus, the characterization must simultaneously describe the *entire* equilibrium correspondences for all initial priors.

We are going to simplify our description of the equilibrium correspondences by focusing on the payoffs of the positive probability types (see, among others, Hart [1985] and Aumann and Hart [2003] for an analogous approach). For any belief system $\pi \in \Pi$, any two payoff vectors $v, v' \in R^{\Theta^*}$, write $v \leq_{\pi} v'$ if $v_i(\theta_i) \leq v'_i(\theta_i)$ for each player *i* type θ_i and $v_i(\theta_i) = v'_i(\theta_i)$ for all π -positive probability types θ_i . In other words, vector v' contains exactly the same payoffs for π -positive probability types and possibly higher payoffs for the zero-probability types. For each set $X \subseteq R^{\Theta^*}$, define

$$X^{\pi+} = \{ v' : \exists v \in A \text{ st. } v \preceq_{\pi} v' \}.$$

For any payoff correspondence $E(\pi)$, define the enhancement of E as a payoff correspondence E^+ such that $E^+(\pi) = (E(\pi))^{\pi+}$ for each π . If $E = E^+$, we say that correspondence E is *enhanced*.

[&]quot;uniform" convergence across many histories. (Notice that the original definition of sequential equilibrium from Kreps and Wilson [1982b] applies only to finite games and the above issue does not arise.)

2.5. Comments. Our analysis of repeated games of incomplete information is restricted to the known-own payoffs case. One of the difficulties in extending the result to the general case is connected with the design of punishment strategies. In the known-own payoffs case, the set of *i*'s individually rational payoffs, as well as -i's strategy that minmaxes player *i*, does not depend on the beliefs of player *i* or the type of players -i (see section 2.3 above or Lemma 4 in Appendix B). This fact allows us to construct equilibria in which minmax strategies are used without any information revelation; we need only to make sure to choose continuation payoffs so that all types of minmaxing players have incentives to randomize with the same probabilities among all pure strategies in the support of the (possibly mixed) minmax strategy.

However, if player *i*'s payoffs depend on the information of player -i, the value of player *i* minmax may depend on his beliefs about the type of player -i. It follows that to punish player *i*, players -i's strategy may depend on -i's types. This complicates the use of punishment in sequential equilibria.

We assume the belief systems have common rectangular support. This restriction ensures that all players can agree with each other (and with an outside observer) on which types have positive or zero probability. Moreover, because Bayesian updating respects the restriction, the restriction is inherited along the equilibrium path. The distinction between the zero and positive probability types is important because their payoffs are analyzed differently (see the above definition of enhanced payoffs). Furthermore, without rectangular support, our notion of individual rationality is not sufficient (see Hörner et al. [2011]).

3. Payoff correspondence

In this section, we construct a set of payoffs in profiles in which (i) information is revealed (i.e., types fully or partially separate) in at most finitely many periods, (ii) in all other periods, types pool their actions, and (iii) at each period, continuation payoffs are individually rational. The construction is inductive. We start with describing payoffs in non-revealing, individually rational profiles. Then we present two inductive steps that correspond to either pooling or information revelation in the initial periods. Payoffs in non-revealing profiles. For each belief system π , let

$$F_0(\pi) = \operatorname{int} \left(IR \cap V^{\pi+} \right). \tag{3.1}$$

It is well known that correspondence F_0 is equal to the payoffs in equilibria in which no information is revealed (see Hart [1985], Koren [1992], and Shalev [1994] for Nash equilibrium and no-discounting, and Peski [2008], and Hörner and Lovo [2009] for the sequential equilibrium in the discounted case). To get some intuition, suppose that π is a full-support belief system. Then, $F_0(\pi) = \operatorname{int}(IR \cap V)$. If such a set is non-empty, there exists a (possibly correlated) action profile a such that $g(a) \in$ $F_0(\pi)$. One can construct equilibria in which a is played on the equilibrium path and deviations by a single player are punished with (Blackwell) minmaxing.

In some games, correspondence F_0 might be empty-valued for some prior belief systems π . Such games do not have non-revealing equilibria.⁸

Initially pooling actions. For each correspondence F, define correspondence $\mathcal{A}F$: for each belief π , let

$$\mathcal{A}F(\pi) = \operatorname{int}\left(IR \cap \operatorname{con}\left(F(\pi) \cup V\right)\right).$$

Correspondence $\mathcal{A}F$ contains all payoffs u that are individually rational and that can be obtained as the convex combination $u = \beta v + (1 - \beta) u'$ of a (possibly, not individually rational) non-revealing payoff $v \in V$ and payoff vector $u' \in F(\pi)$. If u'is an expected payoff in some strategy profile, then u is a payoff in a profile in which, initially, the players' types pool their behavior on profile a. After t periods, where $\delta^t \approx \beta$, players continue with the original profile with payoffs u'.

Revelation of information. Information is revealed (possibly only partially) whenever different types play different (possibly mixed) actions.

⁸For a simple example, consider a game, in which player 1 has two types $\theta_1 \in \{0, 1\}$, two actions $a_1 \in \{0, 1\}$, and receives payoff 1 if his action is equal to his type and 0 otherwise. Then for any prior π that assigns positive probability to both types of player 1, the sets IR and $V^{\pi+} = V$ are disjoint. Intuitively, in any equilibrium, player 1 will always match his action to the state, which immediately reveals all information.

Let $\pi \in \Pi$ be the initial belief system. We represent the revelation of information in the form of *continuation lottery* $l = (\alpha, u)$, where $\alpha = (\alpha_i)$ is a profile of the firstperiod strategies $\alpha_i : \Theta_i \to \Delta A_i$, and $u : A \to R^{\Theta^*}$ is an assignment of continuation payoffs following the realization of the first-period actions. Each lottery must satisfy two conditions. First, we assume that if an action a_i is played with positive probability by some type (i.e., there exists θ_i such that $\alpha_i (a_i | \theta_i) > 0$), then a_i is played with positive probability by some π -positive probability type. This allows us to use the Bayes formula to compute the posterior belief system $p^{\pi,l}(a) = \left(p^{\pi,l,\theta_i}(a_{-i})\right)_{i,\theta_i}$ following positive probability action profile a. Notice that the beliefs of player i depend only on the actions chosen by other players. Second, we require that lottery l ensures that all types of all players are indifferent among all positive probability actions:

$$E_{\pi^{\theta_i}} u_i \left(a_i, \alpha_{-i} \left(\theta_{-i} \right), \theta_i \right) = E_{\pi^{\theta_i}} u_i \left(\alpha_i \left(\theta_i \right), \alpha_{-i} \left(\theta_{-i} \right), \theta_i \right).$$

$$(3.2)$$

(The payoff consequences of playing actions during one period can be ignored.) Define the value of lottery l as a payoff vector $v^{\pi,l} \in \mathbb{R}^{\Theta^*}$ such that for each player i and type $\theta_i, v_i^{\pi,l}(\theta_i)$ is equal to (3.2). Let $L(\pi)$ denote the set of lotteries that satisfy the above conditions.

The incentive condition (3.2) requires that all types of all players are indifferent among all (positive probability) actions, including actions a_i that type θ_i is not supposed to play with positive probability, $\alpha_i (a_i | \theta_i) = 0$. This is stronger than a typical incentive condition, which requires only weak inequality. However, this is without loss of generality: due to the enhancement property, we can always increase the continuation payoffs of type θ_i after action a_i to replace a weak inequality with equality.

For each correspondence F, define correspondence $\mathcal{B}F$: for each belief π , let

$$\mathcal{B}F(\pi) = \left\{ v^{\pi,l} : l \in L(\pi) \text{ and } u(a) \in F\left(p^{\pi,l}(a)\right) \text{ for each pos. prob. } a \right\}.$$

Set $\mathcal{B}F(\pi)$ contains the values of all lotteries with prior belief π and with posteriors and continuation payoffs that belong to correspondence F.

Candidate payoff correspondence. For any two payoff correspondences $F, G \Rightarrow R^{\Theta^*}$, write $F \subseteq G$ if $F(\pi) \subseteq G(\pi)$ for any belief system. The next result follows immediately from the fact that operations \mathcal{A} and \mathcal{B} are monotonic: for any two correspondences, if $F \subseteq G$, then $\mathcal{A}F \subseteq \mathcal{A}G$ and $\mathcal{B}F \subseteq \mathcal{B}G$.

Theorem 1. There exists the smallest correspondence F^* such that $F_0 \subseteq F^*$, $\mathcal{A}F^* \subseteq F^*$, and $\mathcal{B}F^* \subseteq F^*$. Moreover, $F^* = \bigcup_n F_n^B$, where $F_1^A = \mathcal{A}F_0$, and for each $n \ge 1$, $F_n^B = \mathcal{B}F_n^A$ and $F_{n+1}^A = \mathcal{A}F_n^B$.

The Theorem defines correspondence F^* as the smallest correspondence that contains F_0 and that is closed with respect to operations \mathcal{A} and \mathcal{B} . Additionally, the Theorem provides a method of constructing F^* by alternating application of the two operations to initial correspondence F_0 . Each step has a simple geometric characterization. In general, it is not possible to simplify the description by eliminating any of the steps (section 6.2 contains an example of a game and constructions for which all steps are required).

In some games, correspondence F^* might be empty-valued for some prior belief systems π . Trivially, if correspondence $F_0(\pi)$ is non-empty, then $F^*(\pi)$ is not empty as well. If the game has an open thread (see below), then correspondences $\mathcal{BAF}_0 \subseteq F^*$ (but not necessarily F_0 ; see the example from footnote 8) are non-empty valued for all initial priors.

For future reference, notice that correspondences F^* , F_n^A , and F_n^B are enhanced. This follows from the fact that correspondence F_0 is enhanced, and that operations \mathcal{A}, \mathcal{B} , and taking the limit preserve the enhancement property.

4. Finitely revealing payoffs

In this section, we show that correspondence F^* is a lower bound on the set of payoffs obtained in finitely revealing equilibria. For each n, define the limit payoff correspondences⁹

$$FR_{n}(\pi) = \lim \inf_{\delta \to 1} FR_{n}^{\delta}(\pi).$$

Let $FR(\pi) = \bigcup_n FR_n(\pi)$ be the limit set of payoffs in finitely revealing equilibria.

Theorem 2. For each $\pi \in \Pi$, $F^*(\pi) \subseteq FR^+(\pi)$ and for each n, $F_n^B(\pi) \subseteq F_{n+1}^A(\pi) \subseteq FR_n^+(\pi)$.

⁹The infimum limit $\liminf_{\delta \to 1} FR_n^{\delta}(\pi)$ is defined as the set of payoff vectors v such that for each sequence $\delta_n \to 1$, there exists sequence $v_n \to v$ and such that $v_n \in FR_n^{\delta_n}(\pi)$. It is the greatest lower bound on the set of accumulation points.

The proof goes by induction on n and is found in Appendix B. In each step, we construct finitely revealing equilibria with the required payoffs. The constructions are relatively standard and rely on techniques from Fudenberg and Maskin [1986] that are adapted to games with incomplete information. (See also constructions used in Cripps and Thomas [2003] and Peski [2008].)

5. Equilibrium payoffs

In this section, we state our main assumption and show that under this assumption, all Nash equilibrium payoffs are contained in the closure of correspondence F^* .

5.1. **Open thread assumption.** For each type profile $\theta = (\theta_i)_i \in \Theta$, let $\pi^{\theta} \in \Pi$ be the belief system in which all types of player *i* assign probability one to the opponents' profile θ_{-i} .

A thread is an assignment $u^* : \Theta \to R^{\Theta^*}$ of payoff vectors to type profiles such that (a) for each type profile $\theta \in \Theta$, $u^*(\theta)$ is an (enhanced) payoff vector in a non-revealing equilibrium in a game with initial beliefs π^{θ} ,

$$u^*(\theta) \in \mathrm{cl}F_0\left(\pi^{\theta}\right),$$

and (b) for each player *i*, all types $\theta_i, \theta'_i \in \Theta_i$, and all type profiles $\theta_{-i} \in \Theta_{-i}$,

$$u_{i}^{*}\left(heta_{i}, heta_{-i}
ight)=u_{i}^{*}\left(heta_{i}^{\prime}, heta_{-i}
ight)$$
 ,

(Here, $u_i^*(\theta) \in R^{\Theta_i}$ is a vector of payoffs $u^*(\tau_i|\theta)$ for each type τ_i of player *i* given profile $\theta = (\theta_i, \theta_{-i})$ of all types of all players. The above equation is an equality between two vectors.) We say that there exists an *open thread* if u^* can be chosen so that $u_i^*(\theta) \in F_0(\pi^{\theta})$.

A thread has a natural interpretation. Consider a direct revelation mechanism in which players report their types, and following report $\tilde{\theta}$, each player *i* type θ_i receives payoff $u_i^*(\theta_i|\tilde{\theta})$. The first condition ensures that the payoffs can be approximated by equilibrium payoffs in the game with "complete information" (i.e., in which all information is revealed). The second condition ensures *ex post* incentive compatibility for player *i*: If we were to interpret θ_i as a report of player *i*, then player *i* would be indifferent between reporting her type truthfully and reporting any other type *regardless* of the reports of the other player. For each open thread u^* , each $\pi \in \Pi$, define $u^*(\pi) \in \mathbb{R}^{\Theta^*}$ so that for each player i type θ_i ,

$$u_{i}^{*}\left(\theta_{i}|\pi\right)=\sum_{\theta_{-i}}\pi^{\theta_{i}}\left(\theta_{-i}\right)u^{*}\left(\theta_{i}|.,\theta_{-i}\right),$$

 $(u^*(\theta_i|, \theta_{-i}))$ is equal to θ_i -coordinate of the payoff vector $u^*(\theta'_i, \theta_{-i})$ for some θ'_i ; by assumption, this value does not depend on the choice of θ'_i . It follows directly from the definitions that $u^*(\pi) \in F_1^B(\pi)$ for each π . By Theorem 2, $u^*(\pi)$ is a payoff vector in a fully and immediately revealing equilibrium of the game with initial beliefs π and sufficiently high δ . We say that $u^*(\pi)$ forms a multi-linear thread of equilibrium payoffs that passes through games $\Gamma(\pi, \delta)$ for each $\pi \in \Pi$.

All games with two players and one-sided incomplete information have a thread. This follows from the analysis of the non-discounted games in Shalev [1994] (see also Peski [2008]). Additionally, in the case of two players, the threads are essentially equivalent to payoffs in belief-free equilibria of Hörner and Lovo [2009]. The existence of a thread is a necessary condition, and of an open thread is a sufficient condition for the existence of such equilibria (see Appendix A).

5.2. Main result. Our main result provides a characterization of the set of equilibrium payoffs. Define the limit payoff correspondence¹⁰

$$NE(\pi) = \limsup_{\delta \to 1} NE^{\delta}(\pi).$$

Theorem 3. If there exists an open thread, then

$$clNE^{+}(\pi) = clF^{*}(\pi)$$
.

The Theorem provides a characterization of the limit set of payoffs in Nash equilibria of repeated games with incomplete information. The Theorem shows that all Nash equilibrium payoffs for a sufficiently high discount factor can be approximated by payoffs in finitely revealing equilibria that were constructed in Theorem 2.

The Theorem extends a standard folk theorem from the repeated games with complete information. (Note that the existence of the thread is trivial with complete

¹⁰The supremum limit $\limsup_{\delta \to 1} NE^{\delta}(\pi)$ is defined as the set of payoff vectors v such that there exist sequences $v_n \to v$ and $\delta_n \to 1$, such that $v_n \in NE_n^{\delta_n}(\pi)$. It is the smallest upper bound on the set of accumulation points.

information, and that the open thread ensures the appropriate dimensionality conditions are satisfied.) Because finitely revealing equilibria are sequential, the Theorem shows that repeated games with incomplete information preserve a folk-theorem-like feature of games with complete information in which all payoffs in Nash equilibria can be approximated by payoffs in subgame perfect equilibria.

Together with Theorem 1, Theorem 3 provides a method for computing the set of equilibria payoffs. We illustrate the method with examples in section 6.

We explain below that the open thread assumption plays an important role in the proof. We do not know whether the result holds in games that do not satisfy the assumption.

The proof shows that any Nash equilibrium profile in game $\Gamma(\pi, \delta)$ with expected payoffs v can be modified into a profile with expected payoffs that belong to $F^*(\pi)$ and that are arbitrarily close to v. The idea is to modify the original Nash profile to pull the continuation payoffs toward the multi-linear thread u^* . Once the continuation payoffs get sufficiently close to the thread, we conclude the modified profile with one period of full revelation of information, followed by an equilibrium of the "complete" information game.

To get some intuition, suppose that v is a payoff in a Nash profile σ in which during the first period the players choose non-revealing action profile a (i.e., all types of each player i play the same action a_i). Let v(a) be the equilibrium continuation payoffs (we can always choose strategies in such a way so that the continuation payoff after positive probability history is a payoff in a Nash equilibrium). Then v is a convex combination of instantaneous payoffs g(a) and equilibrium continuation payoffs v(a), $v = (1 - \delta) g(a) + \delta v(a)$. See Figure 5.1.

Suppose that v' is a payoff vector that is a convex combination between v and the value of the thread $u^*(\pi)$, $v' = \gamma v + (1 - \gamma) u^*(\pi)$. We can find vector v'(a) (and γ' and δ') such that

• $v' = (1 - \delta') g(a) + \delta' v'(a)$ is a convex combination between v'(a) and g(a). Thus, we can interpret v' as a payoff in a profile that starts with action a, followed by continuation payoffs v'(a), in a game with discount factor $\delta' > \delta$.





• $v'(a) = \gamma' v(a) + (1 - \gamma') u^*(\pi)$ is a convex combination between v(a) and the thread $u^*(\pi)$.

Simple algebra shows that

$$\gamma = \frac{\gamma'}{\gamma'\left(1-\delta\right)+\delta}$$

which implies that $\gamma' < \gamma$. Thus, the relative distance between v'(a) and the thread $u^*(\pi)$ is smaller than the relative distance between v' and the thread.

The above argument applies to situations in which no information is revealed in the first period. If some information is revealed, we show that the relative distance of the (modified) continuation payoffs to the value of the thread in games with new posterior beliefs is smaller than the relative distance of the (modified) payoffs to the thread in the game with prior beliefs. The argument relies on the fact that the expected payoff in the continuation lottery is a convex combination of payoff vectors

$$u(a) = (1 - \delta) g(a) + \delta v(a)$$

with weights on u(a) equal to the prior probability of profile a, and that the prior belief is a convex combination of the posterior beliefs p(a) with exactly the same

weights. The multi-linearity of the thread u^* is essential for the argument (and it is the only place where this assumption is used).

Formally, Theorem 3 follows from two inclusions

$$F^*(\pi) \subseteq FR^+(\pi) \subseteq NE^+(\pi) \tag{5.1}$$

and

$$NE^{+}(\pi) \subseteq \operatorname{cl} F^{*}(\pi) \,. \tag{5.2}$$

The first inclusion is a consequence of Theorem 2. We need to show the other inclusion.

Suppose that $u^*(\pi)$ is an open thread. Choose r > 0 so that for all type profiles θ ,

$$B\left(u^*\left(\pi^{\theta}\right),r\right)\subseteq F_0\left(\pi^{\theta}\right).$$

For each $\delta < 1$, define $\gamma_1^{\delta} = \frac{r}{2M}$. (Note that γ_1^{δ} does not depend on δ .) For each n > 1, inductively define

$$\gamma_n^{\delta} = \frac{\gamma_{n-1}^{\delta}}{\gamma_{n-1}^{\delta} \left(1-\delta\right) + \delta} \in \left(\gamma_{n-1}^{\delta}, 1\right).$$
(5.3)

Notice that $\gamma_n^{\delta} > \gamma_{n-1}^{\delta}$ and $\lim_{n \to \infty} \gamma_n^{\delta} = 1$. Inclusion (5.2) follows from the following result.

Lemma 1. For each n such that $(1 - \gamma_n^{\delta})r > (1 - \delta)M$, for each $\pi \in \Pi$, each $v \in NE^{\delta}(\pi)$,

$$\gamma_{n}^{\delta}v + \left(1 - \gamma_{n}^{\delta}\right)u^{*}\left(\pi\right) \subseteq intF_{n}^{B}\left(\pi\right).$$

5.3. **Proof of Lemma 1.** The proof of Lemma 1 goes by induction on n. First, we show the inductive claim for n = 1. Because $||v|| \leq M$ for each $v \in NE^{\delta}(\pi)$, we have

$$\frac{r}{2M}v + \left(1 - \frac{r}{2M}\right)u^*\left(\pi\right) \in B\left(u^*\left(\pi\right), r\right) \subseteq F_1^B\left(\pi\right).$$

The inclusion comes from Theorem 2 and the definition of an open thread.

Next, suppose that the inductive claim holds for n-1. Take any prior beliefs π and Nash payoff vector $v \in NE^{\delta}(\pi)$. Find an equilibrium profile σ that supports v. Say that action a_i is played with positive probability by player i in the first period if there exists π -positive probability type θ_i such that $\sigma_i(a_i|\emptyset, \theta_i) > 0$. Let A_i^0 denote the set of actions played with positive probability by i.

20

We assume without loss of generality that the continuation strategies are the best responses for all players and all types after all positive probability histories. (If Nash profile σ does not have such a property, the profile can be easily modified without affecting the initial payoffs and equilibrium conditions.)

Non-revealing payoffs. For each positive probability action profile $a \in A^0 := \times_i A_i^0$, each type θ_i , let

$$v\left(a\right) = \left(v_{i}^{p\left(a\right)}\left(\sigma\left(a,.\right);\theta_{i}\right)\right)_{i,\theta_{i}} \in R^{\Theta^{*}}$$

be the vector of continuation payoffs after a. Because a occurs with positive probability, v(a) is a Nash equilibrium payoff in game $\Gamma(p(a), \delta)$. By the inductive assumption,

$$\gamma_{n-1}^{\delta}v\left(a\right) + \left(1 - \gamma_{n-1}^{\delta}\right)u^{*}\left(p\left(a\right)\right) \in \operatorname{int} F_{n-1}^{B}\left(p\left(a\right)\right).$$

Define

$$u(a) = (1 - \delta) g(a) + \delta v(a).$$

Using (5.3), we get

$$\gamma_n^{\delta} u\left(a\right) + \left(1 - \gamma_n^{\delta}\right) u^*\left(p\left(a\right)\right)$$
$$= \gamma_n^{\delta} \left[\delta v\left(a\right) + \left(1 - \delta\right) g\left(a\right)\right] + \left(1 - \gamma_n^{\delta}\right) u^*\left(p\left(a\right)\right)$$
(5.4)

$$= \left(1 - (1 - \delta)\gamma_n^{\delta}\right) \left[\gamma_{n-1}^{\delta}v\left(a\right) + \left(1 - \gamma_{n-1}^{\delta}\right)u^*\left(p\left(a\right)\right)\right] + (1 - \delta)\gamma_n^{\delta}g\left(a\right)$$
(5.5)

$$\in \operatorname{intcon}\left(F_{n-1}^B\left(p\left(a\right)\right) \cup V\right).$$

Because v(a) is a payoff in a Nash equilibrium, $v(a) \in IR$. Because $(1 - \delta) M \leq (1 - \gamma_{n-1}^{\delta}) r$, it must be that

$$\gamma_n^{\delta} \left[\delta v \left(a \right) + \left(1 - \delta \right) g \left(a \right) \right] + \left(1 - \gamma_n^{\delta} \right) u^* \left(p \left(a \right) \right) \in \text{int} IR.$$
(5.6)

Then (5.5) and (5.6) imply that for each positive probability a,

$$\gamma_n^{\delta} u\left(a\right) + \left(1 - \gamma_n^{\delta}\right) u^*\left(p\left(a\right)\right) \in F_n^A\left(p\left(a\right)\right).$$
(5.7)

Revelation of information. For each π -positive probability type θ_i , let

$$\alpha_i\left(\theta_i\right) = \sigma_i\left(\varnothing, \theta_i\right) \in \Delta A_i^0.$$

For each π -zero probability type θ_i , let

$$\alpha_{i}(\theta_{i}) \in \arg \max_{a_{i} \in A_{i}^{0}} u(a_{i}, \alpha_{-i}).$$

Because profile σ is a Nash equilibrium and because of the choice of $\alpha_i(\theta_i)$, for all types θ_i , all positive probability actions a_i ,

$$E_{\pi^{\theta_{i}}}u_{i}\left(a_{i},\alpha_{-i}\left(\theta_{-i}\right),\theta_{i}\right) \leq E_{\pi^{\theta_{i}}}u_{i}\left(\alpha_{i}\left(\theta_{i}\right),\alpha_{-i}\left(\theta_{-i}\right),\theta_{i}\right).$$

The inequality turns into equality for all actions a_i that are played with positive probability by type θ_i . We can replace the inequality with equality for all actions a_i by enhancing the continuation payoffs $u(\theta_i|a)$ of types θ_i that do not play action a_i in strategy α_i (i.e., $\alpha_i(a_i|\theta_i) = 0$). Because correspondence F_n^A is enhanced (see the remark at the end of section 3), (5.7) holds for the enhanced continuation payoffs.

The above implies that the continuation lottery $l = (\alpha, u)$ satisfies (3.2) and belongs to the set $L(\pi)$ (we use the same symbol u to denote the enhanced continuation payoffs). Consider lottery $l' = (\alpha, \gamma_n^{\delta} u(.) + (1 - \gamma_n^{\delta}) u^*(p(.)))$. The properties of the thread u^* imply that for each positive probability $a_i \in A_i$, all types θ_i, θ_{-i} ,

$$E_{\pi^{\theta_{i}}}E_{\alpha_{-i}(\theta_{-i})}u^{*}\left(\theta_{i}|p\left(a_{i},a_{-i}\right)\right)$$

$$=\sum_{\theta_{-i},a_{-i},\theta_{-i}'}\pi^{\theta_{i}}\left(\theta_{-i}\right)\alpha_{-i}\left(a_{-i};\theta_{-i}\right)p^{\theta_{i}}\left(\theta_{-i}'|a_{i},a_{-i}\right)u_{i}^{*}\left(\theta_{i}|.,\theta_{-i}'\right)$$

$$=\sum_{\theta_{-i}'}\pi^{\theta_{i}}\left(\theta_{-i}'\right)u_{i}^{*}\left(\theta_{i}|.,\theta_{-i}'\right)$$

$$=u_{i}^{*}\left(\theta_{i}|\pi\right).$$
(5.8)

In particular, the first line of (5.8) does not depend on positive probability action a_i . Together with the fact that lottery $l \in L(\pi)$, the above implies that lottery l' satisfies (3.2) for each type θ_i .

The value of lottery l' is equal to

$$v^{\pi,l'} = \gamma_n^{\delta} v^{\pi,l} + \left(1 - \gamma_n^{\delta}\right) u^*(\pi) \,,$$

where $v_i^{\pi,l}$ is the value of lottery *l*. Then (5.7) implies that

$$\gamma_{n}^{\delta}v^{\pi,l} + \left(1 - \gamma_{n}^{\delta}\right)u^{*}\left(\pi\right) = v^{\pi,l'} \in \operatorname{int} F_{n}^{B}\left(\pi\right).$$

Notice that

$$v_i^{\pi,l}(\theta_i) = v_i(\theta_i)$$
 for π -positive probability θ_i ,
 $v_i^{\pi,l}(\theta_i) \le v_i(\theta_i) = v_i^{\pi,\delta}(\sigma;\theta_i)$ for π -zero probability θ_i .

The latter follows from the fact that action $\alpha_i(\theta_i)$ is not necessarily the best response action of zero probability type θ_i . Because correspondence F_n^B is enhanced,

$$\gamma_n^{\delta} v + \left(1 - \gamma_n^{\delta}\right) u^*(\pi) \in F_n^B(\pi).$$

This ends the proof.

5.4. Quality of approximation. The proof of Theorem 3 leads to the following bounds on the quality of the approximation of the Nash equilibrium set by *n*-revealing sets F_n^B . Recall that M is an upper bound on the absolute value of the payoffs and r > 0 is the size of the open thread.

Corollary 1. Let $A = \max\left\{\frac{2M}{r}, 2\right\}$. For each $v \in NE^{\delta}(\pi)$, each $\epsilon > (1 - \delta)A$, and either $n \ge \left\lceil \frac{2\log 2A}{\epsilon(1-\delta)} \right\rceil$, or $n \ge \frac{1}{(1-\delta)^2}$, $(1 - \epsilon)v + \epsilon u^*(\pi) \in F_n^B(\pi)$.

Proof. We show first that for each $\delta \geq \frac{1}{2}$ and each $\epsilon > 0$, if $n \geq \left\lceil \frac{\log 2A}{\epsilon(1-\delta)} \right\rceil + 1$, then $\gamma_n^{\delta} \geq 1 - \epsilon$. If not, then $\gamma_1^{\delta} \leq \dots \leq \gamma_n^{\delta} \leq 1 - \epsilon$, and

$$\gamma_n^{\delta} \ge \frac{1}{\delta + (1 - \delta)(1 - \epsilon)} \gamma_{n-1}^{\delta} = \frac{1}{1 - (1 - \delta)\epsilon} \gamma_n^{\delta} \ge \left(\frac{1}{1 - (1 - \delta)\epsilon}\right)^{n-1} \frac{1}{2A},$$

where the last inequality follows from the definition of $\gamma_1^{\delta} = \frac{r}{2M}$. Because

$$-\log\left(1-\epsilon\left(1-\delta\right)\right) \ge \epsilon\left(1-\delta\right),$$

we have a contradiction:

$$\gamma_n^{\delta} \ge e^{(n-1)\epsilon(1-\delta)} \frac{1}{2A} \ge 1 > 1 - \epsilon.$$

Fix $v \in NE^{\delta}(\pi)$. Take any $\epsilon > (1 - \delta) A$. By Lemma 1 and the convexity of set $F_n^B(\pi)$, $\gamma v + (1 - \gamma) u^*(\pi) \in F_n^B(\pi)$ b for each $n \ge \left\lceil \frac{\log 2A}{\epsilon(1-\delta)} \right\rceil + 1$ and any $\gamma \le \gamma_n^{\delta}$ such that $1 - \gamma \le 1 - (1 - \delta) A$. Letting $\gamma = 1 - \epsilon$ establishes the first result.

For the second result, take $\epsilon = (1 - \delta) A$, and observe that for $A \ge 2$, $\frac{\log 2A}{A} \le 1$. The result follows from the first part.

6. Examples

Theorem 1 describes an algorithm for finding all the finitely revealing payoffs. In this section, we illustrate the algorithm with two examples. Our first goal is to illustrate that the algorithm is tractable. Additionally, we want to clarify the relationships between different refinements of Nash equilibria.

In the first example, we show that in a wide class of oligopoly games, all Nash equilibrium payoffs can be obtained in 1-revealing equilibria. At the same time, the set of 1-revealing equilibria can be strictly larger than the set of belief-free equilibria. In fact, we construct an oligopoly game in which the first-best payoffs can be obtained in a 1-revealing equilibrium, but not in a belief-free equilibrium.

In the second example, a bargaining game with one-sided incomplete information, the set of equilibrium payoffs is substantially larger than 1-revealing payoffs. In fact, the equilibria that yield the maximal payoff for the uninformed party typically involve a large number of revelation periods. Although all *n*-revealing equilibria are needed to completely describe the set of payoffs, we are able to derive an explicit description of the payoff set using a solution to a certain differential equation.

6.1. **Oligopoly.** We describe an abstract model of competition that encompasses, as special cases, textbook examples of Bertrand and Cournot competitions with undifferentiated products and incomplete information about the costs. We keep the notation from the general model of a repeated game. There are I firms on the same market. We make two sets of assumptions. The first assumption says that each payoff vector can be replicated by a scheme in which each firm spends a fraction of the period as a single firm on the market, while the other firms are inactive. To state it formally, let $M_i \subseteq R^{\Theta_i}$ be the convex hull of the set of payoff vectors attainable by firm i if firm i was the only firm on the market. We refer to M_i as the set of monopoly payoffs. We assume that M_i is compact, that it contains the zero-payoff vector $\mathbf{0}_i \in M_i$, and that the intersection of M_i with the set of strictly positive payoff vectors has a non-empty interior.

We assume that each payoff vector in the game between the firms is a convex combination of monopoly-inactive payoffs: for each $v \in V$, there exist monopoly payoff $m_i \in M_i$ and market share $\beta_i \geq 0$ for each player *i*, such that $\sum_i \beta_i \leq 1$, and the vector of payoffs of player *i* is equal to $v_i = \beta_i m_i$.

Second, we assume that the set of individually rational payoffs is equal to the set of vectors with non-negative coordinates, $IR = \{v : v_i(\theta_i) \ge 0 \text{ for each } i \text{ and } \theta_i\}$. Any game with payoffs that satisfy the above two assumptions is called an *oligopoly game*. The assumptions imply that the set

$$\left\{ \left(\beta_{1}m_{1},...,\beta_{I}m_{I}\right):\forall_{i}\beta_{i}>0,\sum\beta_{i}=1,m_{i}\in M_{i},\text{ and }\forall_{\theta_{i}}m_{i}\left(\theta_{i}\right)>0\right\}$$

is open and non-empty. It is easy to see that the above set is contained in the set of non-revealing payoffs $F_0(\pi)$ for each π . Thus, each oligopoly game has a non-empty and open set of non-revealing equilibrium payoffs, and in particular, each such game has an open thread.

If we interpret θ_i as the cost parameter, and actions as quantities or prices, then the above assumptions are satisfied in various oligopoly models.

Example 1. The firms play a Cournot oligopoly. The firms choose quantities $q_i \ge 0.^{11}$ The payoff of firm *i* with cost type θ_i is equal to $q_i(P(\sum q_j) - \theta_i)$ where P(.) is an inverse demand function. Let

$$M_{i} = \operatorname{con}\left\{\left(q(P(q) - \theta_{i})\right)_{\theta_{i}} : q \ge 0\right\} \subseteq R^{\Theta_{i}}$$

be the set of monopoly payoffs of firm *i*. Then given strategy profile $(q_1, ..., q_I)$, the payoff of each firm *i* is equal to the fraction $\sum_{q_i}^{q_i}$ of the monopoly payoff obtained from producing quantity $\sum q_j$. By choosing quantity 0, each firm can ensure that its payoff is not smaller than 0. Moreover, if $\lim_{q \to \infty} P(q) < \inf \Theta_i$, then by choosing a sufficiently large quantity, firm -i can ensure that the profits of firm *i* are not higher than 0.

It is easy to check that if function P is differentiable and $\sup \Theta_i < P(0) < \infty$, then set M_i is compact, that it contains the zero-payoff vector $\mathbf{0}_i \in M_i$, and that

¹¹In this and the next example, we allow the firms to choose from infinitely many actions. See footnote 4 for a discussion on how the basic model must be extended.

the intersection of M_i with the set of strictly positive payoff vectors has a nonempty interior.

Example 2. Another model is a Bertrand oligopoly with demand D(.). The firms choose prices p_i . The payoff of each firm i is equal to $D(p_i)(p_i - \theta_i)$ if the firm i's price is strictly lower than the price of its competitors, $\frac{1}{k}D(p_i)(p_i - \theta_i)$, if k-1 other firms choose the lowest price, and 0 otherwise. The two assumptions of the oligopoly games are satisfied if, for instance, function D is differentiable.

Theorem 4. For each oligopoly game, each belief system $\pi \in \Pi$,

$$clNE^{+}\left(\pi\right) = clF_{1}^{B}\left(\pi\right).$$

Moreover, for each payoff vector $v, v \in clF_1^B(\pi)$ if and only if for each type profile $\theta = (\theta_i, \theta_{-i})$, each firm *i*, there exist monopoly payoffs $m_i^{\theta_i} \in M_i$ and market shares $\beta_i^{\theta} \geq 0$, such that $\sum_i \beta_i^{\theta} \leq 1$ and the following conditions hold:

- (1) Individual rationality: $m_i^{\theta_i}(\theta_i) \ge 0$ for each player *i* and type θ_i ,
- (2) Incentive compatibility: for all θ_i, θ'_i ,

$$v\left(\theta_{i}\right) \geq \left(\sum_{\theta_{-i}} \pi^{\theta_{i}}\left(\theta_{-i}\right) \beta_{i}^{\left(\theta_{i}^{\prime},\theta_{-i}\right)}\right) m_{i}^{\theta_{i}^{\prime}}\left(\theta_{i}\right)$$

The proof can be found in Appendix C. The Theorem provides a characterization of the set of equilibrium payoffs. In particular, any equilibrium payoff v can be approximated by a payoff in a profile in which firms immediately reveal their costs and if θ is the true type profile, then player *i*'s payoff is equal to $\beta_i^{\theta} m_i^{\theta_i}(\theta_i)$. The first condition ensures that individual rationality is satisfied *ex post*, and the second condition ensures that firms have *interim* incentives to reveal their types truthfully (although the incentives are not necessarily *ex post*).

The above characterization allows us to address questions of productive efficiency in the equilibria of the repeated oligopoly.

"Pooling" result of Athey and Bagwell [2008]. As an application, we perform a test of the robustness of the "pooling" result from Athey and Bagwell [2008]. Athey and Bagwell [2008] analyze a Bertrand model from example 2 with

$$D(p) = \begin{cases} 1, & \text{if } p \le r, \\ 0, & \text{if } p > r, \end{cases}$$

for some reservation price r > 0. They show that for a sufficiently large discount factor, and given the log-concave distribution of cost types, in the (*ex ante*) optimal symmetric equilibrium, all players choose the same price and receive the same market share regardless of their (privately known) costs. In other words, one can sustain the best payoff in equilibrium in which no player ever reveals any information. There is no contradiction between Athey and Bagwell [2008]'s result and Theorem 4.¹² First, their characterization of optimal equilibrium is tight for all sufficiently high $\delta < 1$, whereas ours simply says that any equilibrium payoff can be approximated by fully revealing payoffs. In fact, one can construct equilibria in which players fully reveal their costs in the first period and then proceed to ignore the revealed information. Because revealing information is costly for discount factors strictly smaller than 1, it should be avoided in the optimal equilibrium of Athey and Bagwell [2008].

Nevertheless, the "pooling" claim is not robust to modifications of the demand. Define the monopoly payoff vector that maximizes the payoffs of type θ_i among all monopoly payoffs of player *i*:

$$m_{\theta_i}^* = \arg \max_{m \in M_i} m\left(\theta_i\right).$$

In Athey and Bagwell [2008], the optimal monopoly price is equal to r and does not depend on the player's type. In general Cournot or Bertrand models, if the demand function is differentiable, then the optimal monopoly action depends on the cost type.

Corollary 2. Suppose that the monopoly actions $m_{\theta_i}^*$ are not identical for all types of player *i*. Then for any π that assigns positive probability to all types, for all sufficiently

¹²There are other differences between Athey and Bagwell [2008]'s model and ours. For example, their demand specification does not lead to a nonempty interior, and our result does not apply. However, it applies to "nearby" models in which the demand below price r is not completely inelastic. In addition, Athey and Bagwell [2008] work with the continuum-type model, whereas in this paper, we assume that there are only finitely many types. These differences do not seem to be important for this discussion.

high $\delta < 1$, there is no Pareto-optimal equilibrium in which players' behavior does not depend on type.

Proof. Suppose that v is an efficient payoff in a profile in which, on the equilibrium path, the players' behavior does not depend on the type. Then there exists $\beta_i \geq 0$ and $m_i \in M_i$ such that $\sum_i \beta_i \leq 1$, and $v_i = \beta_i m_i$.

For each player, construct a payoff vector m_i^* such that for each type θ_i , $m_i^*(\theta_i) = \max\left\{m_{\theta_i}^*(\theta_i), m_i(\theta_i)\right\} \ge m_i(\theta_i)$, with some inequalities strict. Define payoff vector v^* such that player i payoffs are equal to $v_i^* = \beta_i m_i^*$. The mechanism-design characterization implies that $v^* \in F_1^B(\pi)$. Because $v^*(\theta_i) \ge v(\theta_i)$, with some inequalities strict, this contradicts the fact that vector v is efficient.

Belief-free vs. fully and immediately revealing equilibria. In the above characterization of equilibrium payoffs, firms have *interim* incentives to reveal their private information (i.e., before they learn the true types of the other player). Next, we show with an example that we cannot improve the incentives to hold *ex post* (i.e., conditionally on each of the type of the other player). In particular, we show that there exist efficient repeated equilibria that are fully and immediately revealing but that are not belief-free.

Consider a symmetric Cournot model with two players and two cost types for each player, $\Theta = \{h, l\}$, where h > l > 0. Let $m^q \in R^{\Theta}$ be the monopoly payoff vector from quantity q and let $q_{\theta} = \arg \max_q m^q(\theta)$ be the optimal monopoly quantity of type θ . The monopoly profits are maximized by the firm with low costs and quantity q_l . We assume that the optimum is strict:

$$m^{q_l}(l) > m^{q_h}(l), m^{q_l}(h).$$
 (6.1)

Additionally, we assume that the payoff of the high cost type from quantity q_l is strictly positive, but much smaller than the maximum payoff attainable by this type:

$$m^{q_h}(h) > 6m^{q_l}(h) > 0.$$
 (6.2)

We are interested in strategies that maximize the *ex ante* expected sum of the payoffs of both firms. Because of (6.1), the first best for interior beliefs is attained if and only if there is complete productive efficiency:

REPEATED GAMES WITH INCOMPLETE INFORMATION

(u, f_1, f_2)	Weak	Tough
Weak	2, 2, 2x	0, 4, 1 + 3x
Tough	4, 0, 0	-2, -2, 1 - 3x

TABLE 1. Payoffs in bargaining game.

- if both firm types are equal to θ , one of the firms is inactive, and the other one produces quantity q^{θ} . In a symmetric equilibrium, the two allocations are chosen with equal probability,
- if firm *i* has low costs and firm -i has high costs, firm *i* is active, and it produces quantity q^l , and firm -i is inactive (and produces 0).

We claim that the first best allocation cannot be attained in a belief-free equilibrium. Indeed, notice at least one firm i must expect strictly positive profits in a state in which both firms report l (in symmetric allocation, both firms must receive strictly positive profiles). Because firm i receives zero profits if it reports l and the other firm reports h, type h of firm i does not have $ex \ post$ incentives to reveal its true type if the other firm has low costs.

However, if $\pi_i(h) = \frac{1}{2}$ for both players *i*, then (6.2) implies that the first best profile satisfies *ex ante* incentive compatibility. In particular, the first best expected payoff can be attained in a repeated game equilibrium in which both types are revealed immediately and then the play approximates the efficient allocation.

6.2. Labor union-firm bargaining. Consider the following class of games parametrized with $x \in [0, 1]$. There are two players, a labor union (U) and a firm (F). The firm can be either a normal type, $\theta_F = 1$, or a strong type, $\theta_F = 2$. Each player chooses between two actions, Weak and Tough. The payoffs are given in Table 1.

- When x = 1, the payoffs of the normal and strong types of the firm are equal, and the firm and the union play a multi-period bargaining model with complete information.
- When x = 0, the union U and the normal type have payoffs as in the complete information game. The strong type has a (repeated-game) dominant action to play T in every period. This is an example of a model of reputation with

equal discount factors for two players. The complete information game has strictly conflicting interest (Schmidt [1993]): the normal type has a commitment action T such that the union's best reply gives the union its minmax payoff of 0. Cripps et al. [2005] show a reputational result for this class of games: for any p < 1, and for δ high enough, all Nash equilibrium payoffs of the union and the normal type are close to (4, 0).

• For intermediate x, the payoff of the strong type is a convex combination between the normal type and the completely strong type of the reputation case x = 0. The techniques of Cripps et al. [2005] do not apply. (In fact, as we show, the reputational result does not hold). On the other hand, the game has an open thread assumption, and we can use Theorem 3 to compute the set of equilibrium payoffs.

The goal of this section is to describe an "upper," Pareto-optimal, part of the equilibrium set (the "lower" part can be described in an analogous way). To simplify the exposition, we assume that $x < \frac{1}{5}$. We begin with developing geometric intuition about the set of equilibria. A more formal description follows.

Intuition. We use $\pi \in [0, 1]$ to denote the probability of the normal type. Because the minmax strategy of each player is T, the set of individually rational payoffs is equal to t

$$IR = \{(u, f_1, f_2) : u \ge 0, f_1 \ge 0, f_2 \ge 1 - 3x\}.$$

To represent the payoff sets graphically, we focus here on the payoffs of player Uand the normal type of player F, given that the strong type of player F receives her minmax payoff. Figure 6.1a describes the payoffs of player U and the normal type of player F in the "complete information" games in which the type of player F is known with probability 1. (Precisely, it is the projection of the payoff set on the two-dimensional plane). The large triangle describes the sets of payoffs when $\pi = 1$. It is equal to the set of payoffs at the intersection of the convex hull of feasible and individually rational payoffs in a standard complete information game between player U and the normal type of F. The small filled triangle is the set of payoffs when $\pi = 0$, i.e., when player U believes that she faces the strong type of player F. In such a case, player U expects player F to play T sufficiently often so that the strong type of player



FIGURE 6.1. Finitely equilibrium payoffs in the union-firm bargaining.

F receives payoff at least equal to her minmax. The only feasible payoffs that are individually rational for player U and at the same time consistent with the strong type minmax are obtained when player U plays W sufficiently often. In equilibrium, the normal type of player F best responds by mimicking the behavior of the strong type and playing (often) action T. The set of equilibrium payoffs when $\pi = 0$ converges to (4, 0) when $x \to 0$. The small red triangle is also equal to the set of non-revealing payoffs $F_0(\pi)$ (see formula (3.1)) for each $\pi \in (0, 1)$. Notice also that for each π , $F_0(\pi) = \mathcal{A}F_0(\pi)$.

Figure 6.1b shows payoffs A in an immediately and fully revealing equilibrium in a game with initial prior $\pi = \frac{1}{2}$. Suppose that in the first period, each type of player F reveals herself to player U by choosing one of two different actions. The continuation payoffs are equal to A_1 if player F reveals herself to be the normal type and A_2 otherwise. The continuation payoffs A_i to the sets from Figure 6.1a, hence, they can be supported by continuation equilibrium strategies. Because the payoffs of the normal type are equal in both cases, she is indifferent between revealing herself truthfully and mimicking the strong type. (A more careful analysis shows that the continuation payoffs of the strong type are equal to 1 - 3x in both cases.) Thus, the expected payoff $A = \pi A_1 + (1 - \pi) A_2$ can be supported in an equilibrium.

The cross-hatched polygon on Figure 6.1b is equal to the set of all payoffs obtained by immediate full revelation, followed by a play of non-revealing equilibrium, $F_1^B(\pi) = \mathcal{BAF}_0(\pi)$ for $\pi = \frac{1}{2}$. Due to the characterization by Shalev [1994], it is equal to the set of all equilibrium payoffs in the repeated game without discounting. Moreover, it is also equal to the set of belief-free equilibria of Hörner and Lovo [2009].

Next, we consider 1-revealing equilibrium payoffs which cannot be obtained in immediately and fully revealing equilibria, and hence do not belong to the belief-free or no-discounting payoff sets. Consider a profile in which both players play action T for some number of periods and then continue with a profile with payoffs in a set from Figure 6.1b. The set of payoffs in such profiles is equal to the convex hull of payoffs from Figure 6.1b and point (-2, -2). Its subset that consists only of individually rational payoffs is equal to $F_2^A\left(\frac{1}{2}\right) = \mathcal{A}F_1^B\left(\frac{1}{2}\right)$ and it is shown as the cross-hatched area on Figure 6.1c.

Figure 6.1d shows the set of equilibrium payoffs obtained in a class of 2-revealing equilibria of the game with initial prior $\pi = \frac{3}{4}$. In these equilibria, first, players play action profile (T, T) for number of periods. Next, there is one period of partial information revelation: the normal type randomizes with equal probability between W and T and the strong type plays action T. If action W is observed, the play moves to a continuation equilibrium in the game with prior $\pi = 1$; otherwise, the game moves to one of the equilibria with payoffs in Figure 6.1c.



FIGURE 6.2. Payoffs of the firm's types

The above constructions can be repeated any number of times. At each step, the set of payoffs grows larger. In the limit, we obtain the set of all finitely revealing equilibrium payoffs, and as a consequence of our main result, the set of all equilibrium payoffs. Below, we show that the limit has a tractable, closed-form description as a solution to a certain differential equation.

Notation. To describe the payoff sets, we need some notation. We write $f = (f_1, f_2) \in \mathbb{R}^2$ to denote the payoffs of the two types of player F, and $v = (u, f) \in \mathbb{R}^3$ to denote the vector of the payoffs of both players. For any $f^a \neq f^b$, let $I(f^a, f^b)$ be the interval on a two-dimensional plane that connects f^a and f^b . For any non-co-linear $v^a, v^b, v^c \in \mathbb{R}^3$, for each $f \in \mathbb{R}^2$, let $H^{v^a, v^b, v^c}(f)$ be the unique value such that $(H^{v^a, v^b, v^c}(f), f)$ belongs to the unique affine hyperplane that passes through points $v^x, x = a, b, c$.

Figure 6.2 illustrates the payoffs of the firm's types. We find $f^* = (f_1^*, 1 - 3x)$ such that $f^* \in I(g_F(W,T), g_F(W,W))$. Find $\hat{f} = (0, \hat{f}_2)$ such that $f \in I(g_F(W,T), g_F(T,T))$. Finally, we find $f^{**} = (f_1^{**}, 1 - 3x)$ so that $H^{g(T,W),g(T,T)}(f^{**}) = 0$.

Define sets $A, B', B'' \subseteq \mathbb{R}^2$,

$$A = \operatorname{con} \{f^{**}, f^{*}, g_{F}(W, T)\},\$$

$$B' = \operatorname{con} \{f^{**}, (0, 1 - 3x), g_{F}(W, T)\},\$$

$$B'' = \operatorname{con} \{\hat{f}, (0, 1 - 3x), g_{F}(W, T)\}.\$$

Sets A, B', and B'' are illustrated in Figure 6.2.

For each $f \in B'$, choose $j'(f) \in [0, f^{**}]$ so that f belongs to the interval $I(g_F(W, T), (j'(f), 1-3x))$ Similarly, for each $f \in B''$, choose $j''(f) \in [1-3x, \hat{f}]$ so that f belongs to the interval $I(g_F(W, T), (0, j''(f)))$.

We say that function u^{π} describes the upper surface of equilibria if for each f,

$$u^{\pi}(f) = \sup \{ u : (u, f) \in F^{*}(\pi) \}$$

(we take $u^{\pi}(f) = -\infty$, if the right-hand side set is empty).

Complete information payoffs. Using Theorem 5, we can describe the "upper" surface of the payoffs in the complete information case $\pi \in \{0, 1\}$. Let

$$u^{1}(f) = \begin{cases} 4 - f_{1}, & \exists f' \in A \cup B' \cup B'' \text{ st. } f' \preceq_{1} f, \\ -\infty, & \text{otherwise.} \end{cases}$$

$$u^{0}(f) = \begin{cases} \min\left\{H^{g(T,W),g(T,T),g(W,T)}(f), 4 - \frac{4}{1+3x}f_{2}\right\}, & \exists f' \in A \text{ st. } f' \preceq_{0} f, \\ -\infty, & \text{otherwise.} \end{cases}$$

Finitely revealing payoffs. We use our characterization to construct the upper surfaces of the equilibrium sets. First, we construct a sequence of payoff vectors v_n that belong to a finitely revealing set in the game with initial belief $p_n = \frac{n}{N}$, where $N < \infty$. Next, we take $N \to \infty$ and show that the constructed path of equilibria converges to the solution of a certain differential equation.

First, consider the game with initial belief $p_0 = 0$. Let $j_0 = f^{**}$. Due to the above description of the upper surfaces in the complete information case, $v_0 = (0, j_0, 1 - 3x) \in F^*(0)$.

Next, consider the game with initial beliefs p_1 . Vector

$$v' = \frac{1 - p_1}{1 - p_0} \left(0, j_0, 1 - 3x \right) + \frac{p_1 - p_0}{1 - p_0} \left(u^1 \left(j_0, 1 - 3x \right), j_0, 1 - 3x \right)$$

is equal to the value of the p_1 -incentive-compatible lottery in which the firm's normal type gets revealed with probability $\frac{p_1-p_0}{1-p_0}$, upon which the players' continuation payoffs are equal to $(u^1(j_0, 1-3x), j_0, 1-3x)$. If the normal type is not revealed, the labor union updates its belief to p_0 , and the play continues with payoffs $(0, j_0, 1-3x)$. Because of stage B of the construction of the finitely revealing set (Lemma 6), $v' \in F^*(p_1)$.

34

Further, construct a profile in which players play actions (T,T) for fraction $\alpha \in (0,1)$ of time, and then continue with a profile that leads to payoffs v'. The payoffs in such a profile are equal to

$$v_1 = \alpha g \left(T, T \right) + \left(1 - \alpha \right) v'.$$

We choose α so that the payoff of the labor union in vector v is equal to 0. Then by stage A (Lemma 7), $v_1 = (0, j_1, 1 - 3x) \in F^*(p_1)$, where

$$j_1 = \frac{(p_1 - p_0) u^1 (j_0, 1 - 3x)}{1 - p_0} \left(2 + \frac{p_1 - p_0}{1 - p_0} u^1 (j_0, 1 - 3x) \right)^{-1} (2 + j_0).$$

Using the same argument, we show that if $v_n = (0, j_n, 1 - 3x) \in F^*(p_n)$, and j_n is not too close to 0, then $v_{n+1} = (0, j_{n+1}, 1 - 3x) \in F^*(p_{n+1})$, where

$$j_{n+1} = j_n + \frac{(p_{n+1} - p_n)u^1(j_n, 1 - 3x)}{1 - p_n} \left(2 + \frac{p_{n+1} - p_n}{1 - p_n}u^1(j_n, 1 - 3x)\right)^{-1} (2 + j_n).$$

After some algebraic transformations, we obtain

$$\frac{p_{n+1} - p_n}{j_{n+1} - j_n} = \frac{2 + \frac{p_{n+1} - p_n}{1 - p_n} u^1 \left(j_n, 1 - 3x\right)}{2 + j_n} \frac{1 - p_n}{u^1 \left(j_n, 1 - 3x\right)}$$

By taking limit $N \to \infty$, the above equation converges to the differential equation

$$\frac{dp}{dj} = -\frac{2}{2+j} \frac{1-p(j)}{u^1(j,1-3x)}.$$
(6.3)

(The minus comes from the fact that $p_{n+1} - p_n = -\frac{1}{N}$.)

Suppose that $p' : [0, f^{**}] \to [0, 1]$ is a solution to (6.3) such that $p'(f^{**}) = 0$. Choose π^* so that $p'(0) = \pi^*$. The above analysis implies that for each $\pi \leq \pi^*$, each $j \in [0, f^{**}]$,

$$(0, j, 1 - 3x) \in F^*(p'(j)).$$

Because set $F^*(p'(j))$ is convex and it contains vector g(W,T), it must be that $(0, f) \in F^*(p'(j'(f)))$ for each $f \in B'$.

Similar equations can be derived for the elements of set B''. Let $p'': [1 - 3x, \hat{f}] \rightarrow [0, 1]$ be a solution to the following differential equation: $p''(1 - 3x) = \pi^*$, and

$$\frac{dp''}{dj} = -\frac{4/3}{\hat{f} - j} \frac{1 - p''(j)}{u^1(0, j)}.$$
(6.4)

Then for each $f \in B''$, we have $(0, f) \in F^*(p''(j''(f)))$.

Proposition 1. The following functions describe the upper surfaces of equilibria:

- $if \pi \leq \pi^*$, let $u^{\pi}(f) = \begin{cases} \pi u^1(f) + (1 - \pi) u^0(f), & f \in A, \\ \frac{\pi - p'(j'(f))}{1 - p'(j'(f))} u^1(f), & f \in B' \text{ and } \pi \geq p'(j'(f)), \\ -\infty, & otherwise. \end{cases}$
- if $\pi > \pi^*$, let

$$u^{\pi}(f) = \begin{cases} \pi u^{1}(f) + (1 - \pi) u^{0}(f), & f \in A, \\ \frac{\pi - p'(j'(f))}{1 - p'(j'(f))} u^{1}(f), & f \in B' \text{ and } \pi \ge p'(j'(f)) \\ \frac{\pi - p''(j''(f))}{1 - p''(j''(f))} u^{1}(f), & f \in B'' \text{ and } \pi \ge p''(j''(f)) \\ -\infty, & otherwise. \end{cases}$$

Proof. The above discussion shows that $(u^{\pi}(f), f) \in F^{*}(\pi)$ for each $f \in \mathbb{R}^{2}$ such that $u^{\pi}(f) > -\infty$. We are left with showing that for each $u > u^{\pi}(f), (u, f) \notin F^{*}(\pi)$.

Define correspondence $F(\pi) \supseteq \{(u, f) : u \le u^{\pi}(f)\}$ for each π . We will show that none of the operations described in section 3 adds any payoffs to correspondence F.

First, notice that $F(\pi) = IR \cap \operatorname{con} \{V \cup F(\pi)\}.$

Second, we are going to show that each π -incentive-compatible lottery with continuation payoffs in the correspondence F(.) has its value in set $F(\pi)$. Indeed, suppose that $l = (\alpha, \psi)$ is such a lottery with value v = (u, f) and continuation payoffs $\psi(a) = (u(a), f(a))$ after positive probability actions a of the firm. Then $f(a) \leq f$ with equality if action a is played with positive probability by the two types of the firm. Moreover, if action a is played with positive probability by only one type, we can use the description of the upper surfaces in the "complete information" games to show that $u^{p(a)}(f) \geq u^{p(a)}(f(a))$.

Consider lottery $l' = (\alpha, \psi')$, where $\psi'(a) = (u^{p(a)}, f)$ for all actions a. Then the description of the upper surface u^{π} implies that

$$u \le \sum_{a} p(a) u^{p(a)} \le u^{\pi}(f) \, ,$$

which, in turn, implies that $(u, f) \in F(\pi)$.

36

Equilibrium behavior. One can use the above analysis to (approximately) predict the dynamics along the equilibria that support payoffs on the upper surfaces. As an example, we describe the equilibrium behavior that induces (approximately) payoff vector $(0, f_1, 1 - 3x)$ in the game with initial beliefs $p'(f_1)$ for some $f_1 \in [0, f^{**}]$. This profile can be described in three phases.

- In the revelation phase, the labor union and the strong type of player F play Tough. The normal type of F plays Tough with a probability that is close to 1, and with a small probability, the normal type plays Weak. The phase ends either because the normal type reveals herself playing Weak, or because the posterior probability of the normal type becomes equal to 0 (i.e, the strong type is revealed). In the former case, the players continue with the "normal type" phase; in the latter, the players continue with the "strong type" phase. The continuation payoff of the normal type f_1 throughout the revelation phase gradually increases with the decreasing posterior probability $p'(f_1)$ of the normal type. The rate with which the normal type chooses Weak is chosen so that the continuation payoff of the labor union is equal to 0 at each moment of the revelation phase.
- In the "normal type" phase, players play the "complete information" game equilibrium with payoffs equal to $(u^1(p'(f_1)), f_1, 1 3x)$, where f_1 is the expected continuation payoff of the normal type at the moment of revelation.
- In the "strong type" phase, players play the equilibrium of the "complete information" game with payoffs $(0, f^{**}, 1 3x)$.

In a similar way, we can describe strategy profiles that induce any other payoff on the upper surface.

6.3. Labor union-firm bargaining with two-sided incomplete information.

For the sake of completeness, we find it worthy to point out that a version of the above model with two-sided incomplete information does not have to have a thread (examples of games with known-own payoffs and no belief-free equilibria are known in the literature - see Koren [1992] and Hörner and Lovo [2009]). Suppose that there are two types of each player, and the payoffs are given in Table 2.

(u_1, u_2, f_1, f_2)	Weak	$T \text{ough} \\ \hline 0, 0, 4, 1 + 3x$	
Weak	2, 2x, 2, 2x		
Tough	4, 1 + 3x, 0, 0	-2, 1 - 3x, -2, 1 - 3x	

TABLE 2. Payoffs in bargaining game.

Lemma 2. Suppose that $x < \frac{3}{100}$. Then the labor union-firm bargaining game with two-sided incomplete information does not have a thread.

The proof of Lemma 2 can be found in Appendix D.

7. Comparison with no-discounting case

We compare our characterization of payoffs with the characterization from Hart [1985] in the case of no-discounting. Hart [1985] considers the general payoffs case and he assumes that there are two players, uninformed U (with one type) and informed I. Let Θ_I be the finite set of the types of the informed player, and let Δ_{Θ_I} be the simplex of beliefs of the uninformed player. Then the correspondence of the Nash equilibrium payoffs can be characterized as the set of initial values $(v_{U,0}, v_{I,0}, p_0) \in R \times R^{\Theta_I} \times \Delta_{\Theta_I}$ of a class of bi-martingales, i.e., stochastic processes that satisfy the following three properties:

- for all odd t, $p_t = p_{t+1}$ and $E(v_{U,t+1}, v_{I,t+1} | \mathcal{F}_t) = (v_{U,t}, v_{I,t})$,
- for all even $t, v_{I,t} = v_{I,t+1}$ and $E(v_{U,t+1}, p_{t+1} | \mathcal{F}_t) = (v_{U,t}, p_t),$
- the limit payoff $(v_{U,\infty}, v_{I,\infty}) = \lim_{t\to\infty} (v_{U,t}, v_{I,t})$ is a payoff in a repeated game with initial prior $p_{\infty} = \lim_{t\to\infty} p_t$ and in which no further substantial information is revealed. In the known-own payoff case, the set of such payoffs is equal to $F_0(p)$ defined in (3.1).

The second and the third properties are equivalent to, respectively, the revelation of information (operation \mathcal{B}) and the no-revealing payoffs F_0 , from our characterization.

The first property convexifies the set of payoffs obtained in the previous steps. (Recall that Hart does not assume public randomization and, instead, uses Aumann-Maschler's jointly controlled lotteries.) It corresponds to operation \mathcal{A} from our characterization with a *key difference*: in the discounted case, the payoffs are additionally

convexified with the set of feasible and non-revealing payoffs V. To compare the first property of bi-martingales and operation \mathcal{A} side by side, let $E_N : \Delta_{\Theta_I} \rightrightarrows R \times R^{\Theta_I}$ denote the equilibrium payoff correspondence in the undiscounted cases. Then Hart's characterization implies that for each $p \in \Delta_{\Theta_I}$,

$$E_{N}(p) = \operatorname{con}(E_{N}(p)) = \operatorname{con}(E_{N}(p) \cup (V \cap IR)).$$

The second equality comes from the fact that $V \cap IR \subseteq F_0(p) \subseteq E_N(p)$. Because for any set E,

$$\operatorname{con}\left(E \cup (V \cap IR)\right) \subseteq \operatorname{con}\left(E \cup V\right) \cap IR,$$

and the inclusion is typically strict, the set of payoffs in the no-discounting case is included, and it is typically smaller than the set of payoffs in the discounted case (for example, see section 6.2 and the discussion of Figure 6.1).

In the known-own payoffs case, Shalev [1994] provides a much simpler characterization of no-discounting equilibrium payoffs. For each $p \in \Delta_{\Theta_I}$, $E_N(p)$ is equal to payoff vectors (v_U, v_I) such that $(v_U, v_I) \in IR$, and for each type $\theta \in \Theta_I$ of the informed player, there exists $v^{\theta} \in V$ so that

$$v_U = \sum_{\theta} p\left(\theta\right) v_U^{\theta},$$

and for each $\theta, \theta' \in \Theta_I$,

$$v_I(\theta) = v_I^{\theta}(\theta) \ge v_I^{\theta'}(\theta)$$

All such payoffs can be obtained by immediate and full revelation of the informed player's type θ , followed by the equilibrium play of a profile that corresponds to payoff vector v^{θ} . It is easy to show that the set of such payoffs is equal to $F_1^B(p)$ (see Cripps and Thomas [2003]), or the set of payoffs obtained in belief-free equilibria (see Hörner and Lovo [2009]. Note that that the latter is true because Shalev [1994] is limited to the one-sided case. In particular, $F_1^B(p)$ is always weakly included in the set of payoffs attained in the equilibria with discounting, and the inclusion is strict in games in which there exist non-trivial *n*-revealing equilibria for n > 1.

There is another important difference between Hart's characterization and our result. In the general payoff case, there are games with Nash payoffs that cannot be approximated by equilibria with a finite and bounded number of revelations (Forges

[1984], Forges [1990]; see also the "four frogs" example of Aumann and Hart [2003]. Mathematically, the result follows from the fact that the di-span of a set might be strictly larger than its di-convex hull Aumann and Hart [1986].) More importantly, the bi-martingale characterization is not constructive, and no known algorithm exists that allows one to find all the payoffs in the general case. In our case, we show that all equilibrium payoffs can be approximated by payoffs in equilibria with a bounded number of rounds of revelations, and the characterization is constructive.

The reason for the difference is not clear. On one hand, the characterization from Shalev [1994] shows that only one round of revelation is necessary in the nodiscounting case with known-own payoffs. This would suggest that, at least in the one-sided case, the difference is due to the known-own payoffs assumption. On the other hand, we do not know whether one can find a version of the "four frogs" example with known-own payoffs and multi-sided incomplete information. (As far as we know, the characterization of payoffs in such a case remains an open problem.)

8. Conclusions

This paper provides a characterization of the equilibrium payoffs in repeated games with incomplete information, with discounting, known-own payoffs, and permanent types. We assume that there exists an open multi-linear thread of payoffs in equilibria during which in the first period of the game, players fully reveal their information (i.e., all types of each player take separating actions), and such that the players are *ex post* indifferent between revealing their type truthfully or reporting any other type (i.e., they are indifferent conditionally on any type of the opponent). The assumption is generically satisfied in games with one-sided incomplete information, as well as some important examples of games with multi-sided incomplete information.

The characterization says that all Nash equilibrium payoffs can be approximated by payoffs in finitely revealing equilibria. The characterization leads to an algorithm for finding the equilibrium set through a sequence of geometric operations. This algorithm can be implemented numerically. In examples, we show the characterization can be used to find the exact description of the equilibrium sets analytically. The characterization cannot be further simplified. Table 3 contains relations between non-revealing equilibrium payoffs

 \subsetneq belief-free equilibrium payoffs

 \subsetneq fully and immediately revealing equilibrium payoffs

 \subsetneq *n*-revealing equilibria for $1 < n < \infty$

 \subsetneq all finitely revealing payoffs

=all Nash equilibrium payoffs

 \subsetneq feasible and individually rational payoffs.

TABLE 3. Relations between different kinds of equilibria.

different kinds of equilibria in repeated games with incomplete information and discounting. The inclusion " $A \subsetneq B$ " means that for all games, the set of payoffs A is weakly included in B and that there is an example of a game such that the inclusion is strict (specifically, oligopoly game from the last part of section 6.1 is an example for the strictness of the second inclusion, and the bargaining game from section 6.2 is an example for all the other inclusions).

Further work is required to build tools that allow for analytical description in general games. For instance, the equilibrium set in the bargaining problem from section 6.2 is described as a solution to a certain ordinary differential equation. This method can be easily generalized to other games with one-sided uncertainty and two types. We suspect that differential equations play an important role in more general settings (with more types or with multi-sided uncertainty), but we do not know how to do it.

Other questions are left unanswered by this paper. Most importantly, we would like to know whether a similar characterization holds for games in which an open thread assumption is not satisfied (see an example at the end of section 6.2 or Hörner and Lovo [2009]). Our current methods do not allow us to form a hypothesis one way or the other. It would be interesting to check whether the current analysis extends

in some way to the case of persistent types.¹³ We leave these questions for future research.

References

- Susan Athey and Kyle Bagwell. Collusion with persistent cost shocks. *Econometrica*, 76(3):493–540, April 2008. ISSN 1468-0262. doi: 10.1111/j.1468-0262.2008.00845.x.
 1, 1, 6.1, 12, 13
- Robert J. Aumann and Sergiu Hart. Bi-convexity and bi-martingales. Israel Journal of Mathematics, 54(2):159–180, June 1986. ISSN 0021-2172, 1565-8511. doi: 10. 1007/BF02764940. 7
- Robert J. Aumann and Sergiu Hart. Long cheap talk. *Econometrica*, 71(6):1619–1660,
 October 2003. ISSN 1468-0262. doi: 10.1111/1468-0262.00465. 1, 2.4, 7
- Robert J. Aumann and Michael Maschler. Repeated Games with Incomplete Information. The MIT Press, May 1995. ISBN 0262011476. 1
- James Bergin. A characterization of sequential equilibrium strategies in infinitely repeated incomplete information games. Journal of Economic Theory, 47(1):51– 65, 1989. 1
- David Blackwell. An analog of the minimax theorem for vector payoffs. Pacific Journal of Mathematics, 6(1):1–8, 1956. ISSN 0030-8730. 2.3, 2.3, B
- Jimmy Chan. On the non-existence of reputation effects in two-person infinitelyrepeated games. Economics Working Paper Archive 441, The Johns Hopkins University,Department of Economics, 2000. 1
- Martin W. Cripps and Jonathan P. Thomas. Reputation and perfection in repeated common interest games. *Games and Economic Behavior*, 18(2):141–158, 1997. 1
- Martin W. Cripps and Jonathan P. Thomas. Some asymptotic results in discounted repeated games of one-sided incomplete information. *Mathematics of Operations Research*, 28(3):433–462, August 2003. ISSN 0364-765X, 1526-5471. doi: 10.1287/ moor.28.3.433.16394. 1, 3, 4, 7

¹³Athey and Bagwell [2008] introduce a model of persistent types. Escobar and Toikka [2012] prove a folk theorem for limit $\delta \to 1$ and fixed rates of transitions. One can consider an alternative limit $\delta \to 1$ when the probability of transitions scales with $1 - \delta$.

- Martin W. Cripps, Eddie Dekel, and Wolfgang Pesendorfer. Reputation with equal discounting in repeated games with strictly conflicting interests. *Journal of Economic Theory*, 121(2):259–272, April 2005. ISSN 00220531. doi: 10.1016/j.jet.2004.04.009. 1, 6.2
- Juan F. Escobar and Juuso Toikka. A Folk Theorem with Markovian Private Information. 2012. 13
- F. Forges. Note on nash equilibria in infinitely repeated games with incomplete information. *International Journal of Game Theory*, 13(3):179–187, 1984. ISSN 0020-7276. doi: 10.1007/BF01769814. 1, 7
- FranA§oise Forges. Equilibria with communication in a job market example. The Quarterly Journal of Economics, 105(2):375–398, May 1990. ISSN 0033-5533, 1531-4650. doi: 10.2307/2937792. 1, 7
- D. Fudenberg and E. Maskin. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, pages 533–554, 1986. 4
- Drew Fudenberg and Yuichi Yamamoto. Repeated games where the payoffs and monitoring structure are unknown. *Econometrica*, 78(5):1673–1710, October 2010. ISSN 1468-0262. doi: 10.3982/ECTA8565. 1
- Drew Fudenberg and Yuichi Yamamoto. Learning from private information in noisy repeated games. Journal of Economic Theory, 146(5):1733–1769, September 2011. ISSN 0022-0531. doi: 10.1016/j.jet.2011.03.003. 1
- Sergiu Hart. Nonzero-sum two-person repeated games with incomplete information. Mathematics of Operations Research, 10(1):117–153, February 1985. ISSN 0364-765X, 1526-5471. doi: 10.1287/moor.10.1.117. 1, 5, 2.4, 3, 7, B
- Johannes Hörner and Stefano Lovo. Belief-free equilibria in games with incomplete information. *Econometrica*, 77(2):453–487, 2009. (document), 1, 5, 3, 5.1, 6.2, 6.3, 7, 8, A, B, B
- Johannes Hörner, Stefano Lovo, and Tristan Tomala. Belief-free equilibria in games with incomplete information: Characterization and existence. *Journal of Economic Theory*, 146(5):1770–1795, September 2011. ISSN 0022-0531. doi: 10.1016/j.jet. 2011.06.001. 1, 2.5

- G. Koren. Two-person repeated games where players know their own payoffs. Technical report, Courant Institute, 1992. 1, 3, 6.3, B
- David M. Kreps and Robert Wilson. Reputation and imperfect information. *Journal* of Economic Theory, 27(2):253–279, 1982a. 1
- David M. Kreps and Robert Wilson. Sequential equilibria. *Econometrica*, 50(4): 863–94, 1982b. 7
- Paul Milgrom and John Roberts. Predation, reputation, and entry deterrence. Journal of Economic Theory, 27(2):280–312, 1982. 1
- Marcin Peski. Repeated games with incomplete information on one side. *Theoretical Economics*, 3(1), 2008. 1, 1, 5, 3, 4, 5.1, B, B
- Klaus M. Schmidt. Reputation and equilibrium characterization in repeated games with conflicting interests. *Econometrica*, 61(2):325–51, 1993. 1, 6.2
- Jonathan Shalev. Nonzero-sum two-person repeated games with incomplete information and known-own payoffs. Games and Economic Behavior, 7(2):246–259, 1994. 1, 3, 5.1, 6.2, 7, B
- Thomas Wiseman. A partial folk theorem for games with unknown payoff distributions. *Econometrica*, 73(2):629–645, February 2005. ISSN 1468-0262. doi: 10.1111/j.1468-0262.2005.00589.x. 1

Appendix A. Threads and belief-free equilibria with two players

Hörner and Lovo [2009] give two necessary conditions for the existence of belief-free equilibria in the case of two players. We restate the conditions in our notation and in the known-payoff case. For each probability distribution $\alpha \in \Delta A$, let $g(\alpha) \in V$ be the expectation of payoff vectors $g(\alpha)$ taken with respect to α . Take a pair of vectors $v_i \in R^{\Theta_1 \times \Theta_2}$ for each player i = 1, 2.

- Vectors v_1 and v_2 satisfy *Individual Rationality* if for each player *i*, each type θ_{-i} , the payoffs of player *i* types are individually rational: $\forall \phi \in R^{d_i}_+, \phi \cdot v_i^{,\theta_{-i}} \geq m_i(\phi)$, where $m_i(\phi)$ is the value of the ϕ -weighted minmax defined in (2.1).
- Vectors v_1 and v_2 satisfy *Incentive Compatibility* if for each type profile (θ_1, θ_2) , there exists $\alpha_{\theta_1, \theta_2} \in \Delta A$ such that for each type profile (θ_1, θ_2) , player *i*, type

 $\theta_i',$

$$v_i^{\theta_i,\theta_{-i}} = g_i\left(\alpha_{\theta_i,\theta_{-i}}|\theta_i\right) \ge g_i\left(\alpha_{\theta'_i,\theta_{-i}}|\theta_i\right).$$

The next result shows that the threads are essentially equivalent to the payoff vectors that are Individually Rational and Incentive-Compatible.

Lemma 3. Suppose that $u^* : \Theta_1 \times \Theta_2 \to R^{\Theta^*}$ is a thread. Let v_1 and v_2 be a pair of vectors $v_i \in R^{\Theta_1 \times \Theta_2}$ such that $v_i^{\theta_i, \theta_{-i}} = u^*(\theta_i, \theta_{-i}|\theta_i)$ for each player *i*. Then v_1 and v_2 satisfy Individual Rationality and Incentive Compatibility.

Conversely, suppose that a pair of vectors v_1 and v_2 satisfies Individual Rationality and Incentive Compatibility. For each player i types $\theta_i, \theta'_i \in \Theta_i$, and $\theta_{-i} \in \Theta_{-i}$, let

$$u^*\left(\theta_i', \theta_{-i} | \theta_i\right) = v_i^{\theta_i, \theta_{-i}}$$

Then u^* is a thread.

Proof. Part I. Suppose that u^* is a thread. By the definition of sets $NE(\theta_1, \theta_2)$ from Theorem 5, there exist probability distributions $\alpha_{\theta_i,\theta_{-i}}^{\theta_1^*,\theta_2^*} \in \Delta A$ such that for each type profile (θ_1^*, θ_2^*) , and for each θ_i, θ_{-i} ,

$$u^*\left(\theta_i|\theta_1^*,\theta_2^*\right) = g_i\left(\alpha_{\theta_i,\theta_{-i}^*}^{\theta_1^*,\theta_2^*}|\theta_i\right),$$

and for each player *i* and all types θ_i, θ'_i ,

$$g_i\left(\alpha_{\theta_i,\theta_{-i}^*}^{\theta_1^*,\theta_2^*}|\theta_i\right) \ge g_i\left(\alpha_{\theta_i',\theta_{-i}^*}^{\theta_1^*,\theta_2^*}|\theta_i\right)$$

Define

$$v_i^{\theta_1,\theta_2} = u^* \left(\theta_i, \theta_{-i} | \theta_i \right).$$

Because u^* is a thread, for each player *i*, type θ_{-i} , each type θ'_i

$$v_i^{\theta_i,\theta_{-i}} = u^* \left(\theta'_i, \theta_{-i} | \theta_i \right).$$

Because $u^*(\theta'_i, \theta_{-i}) \in IR$, the payoffs of types of player *i* in the vector $u^*(\theta'_i, \theta_{-i})$ are individually rational. This shows that vectors (v_1, v_2) satisfy Individual Rationality.

Next, we show that (v_1, v_2) satisfies Incentive Compatibility. For each type profile (θ_1, θ_2) , define

$$\alpha_{\theta_1,\theta_2}^* = \alpha_{\theta_1,\theta_2}^{\theta_1,\theta_2} \in \Delta A.$$

Then

$$v_i^{\theta_1,\theta_2} = g\left(\alpha_{\theta_1,\theta_2}^*|\theta_i\right),\,$$

and

$$v_i^{\theta_1,\theta_2} = g\left(\alpha_{\theta_1,\theta_2}^*|\theta_i\right) = g\left(\alpha_{\theta_1,\theta_2}^{\theta_1,\theta_2}|\theta_i\right) = u^*\left(\theta_i|\theta_1,\theta_2\right) = u^*\left(\theta_i|\theta_1',\theta_2\right)$$
$$= g\left(\alpha_{\theta_1,\theta_2}^{\theta_1',\theta_2}|\theta_i\right) \ge g\left(\alpha_{\theta_1',\theta_2}^{\theta_1',\theta_2}|\theta_i\right) = g_i\left(\alpha_{\theta_i',\theta_{-i}}^*|\theta_i\right).$$

Part II. Suppose that the pair of vectors $v_i \in R^{\Theta_1 \times \Theta_2}$ satisfies Individual Rationality and Incentive Compatibility. Let $\alpha_{\theta_1,\theta_2} \in \Delta A$ be as in the definition of Incentive Compatibility. For each profile (θ_1, θ_2) , each player *i* type θ'_i , define

$$u^*\left(\theta_1, \theta_2 | \theta_i'\right) = v_i^{\theta_i', \theta_{-i}} = g_i\left(\alpha_{\theta_i', \theta_{-i}} | \theta_i'\right).$$

Then for each profile (θ_1^*, θ_2^*) , the vector of the payoffs of player *i* types, $u_i^*(.|\theta_1^*, \theta_2^*) = v_i^{.,\theta_{-i}}$, is individually rational. Thus, $u^*(\theta_1^*, \theta_2^*) \in IR$. Moreover, for each profile (θ_1^*, θ_2^*) , and any two types θ_i, θ_i' ,

$$u^*\left(\theta_1^*, \theta_2^* | \theta_i\right) = g_i\left(\alpha_{\theta_i, \theta_{-i}^*} | \theta_i\right) \ge g_i\left(\alpha_{\theta_i', \theta_{-i}^*} | \theta_i\right)$$

This shows that $u^*(\theta_1^*, \theta_2^*) \in NE(\theta_1^*, \theta_2^*)$.

Appendix B. Proof of Theorem 2

The proof of Theorem 2 follows from Lemmas 5, 6, and 7 below. We begin with the preliminary result.

Lemma 4. For each $\epsilon > 0$, there exist $\delta^{\epsilon} < 1$ and $m^{\epsilon} < \infty$ such that for each player *i*, each $m \ge m^{\epsilon}$, and each *v* such that $B(v, \epsilon) \subseteq IR$, there exist *m*-period strategies of players $j \ne i$, $\mu_j^{i,v,m,\epsilon} : \bigcup_{s < m^{\epsilon}} (A_i)^{s-1} \to \Delta A_j$ such that for any sequence $\hat{a}^i = (a_0^i, ..., a_{m^{\epsilon}-1}^i)$ of actions of player *i*, each type θ_i , each $\delta \ge \delta^{\epsilon}$, the following inequality is satisfied:

$$M_{i}^{v,m,\epsilon,\delta}\left(\hat{a}^{i};\theta_{i}\right) := \frac{1-\delta}{1-\delta^{m}} \sum_{s=0}^{m-1} \delta^{s} Eg_{i}\left(a_{s}^{i},\mu_{-i}^{i,v,m,\epsilon}\left(a_{0}^{i},...,a_{s-1}^{i}\right);\theta_{i}\right)$$
$$\leq v_{i}\left(\theta_{i}\right).$$

Here, the expectation is taken over actions induced by strategies $\mu_{-i}^{i,v,m,\epsilon}$.

46

Proof. The Lemma is a discounted version of the Blackwell approachability argument (Blackwell [1956]) (see also Peski [2008] or Hörner and Lovo [2009] for games with discounting). The proof follows the same line and an observation that when $\delta \rightarrow 1$, the discounted payoff criterion in a game with finitely many periods converges to the average payoff criterion.

Lemma 5. For each $\pi \in \Pi$,

$$F_0\left(\pi\right) \subseteq FR_0^+\left(\pi\right).$$

We omit the formal proof, because this result is well-known (see Hart [1985], Koren [1992], and Shalev [1994] for the Nash equilibrium and no-discounting, and Peski [2008], and Hörner and Lovo [2009] for the sequential equilibrium in the discounted case).

Lemma 6. If $F_{n-1}^{B}(\pi) \subseteq FR_{n-1}^{+}(\pi)$, then $F_{n}^{A}(\pi) \subseteq FR_{n-1}^{+}(\pi)$.

Proof. Take any $v^* \in F_n^A(\pi) = \operatorname{int} IR \cap \operatorname{con} \left\{ \operatorname{int} F_{n-1}^B(\pi) \cup V \right\}$. Find $\alpha^* \in (0, 1)$, $g^* \in V$, and $u^* \in \operatorname{conint} F_{n-1}^B(\pi)$ such that $v^* = \alpha^* g^* + (1 - \alpha^*) u^*$. Assume that there exists a pure action profile a^* such that $g(a^*) = g^*$. The assumption is without loss of generality due to public correlation.

Find a sequence of t_{δ} such that $\delta^{t_{\delta}} \to 1 - \alpha^*$ as $\delta \to 1$. We are going to compute the payoffs in a profile in which players play action profile a^* during the initial t^{δ} periods, and then receive continuation payoffs u chosen so that $v^* = (1 - \delta^{t^{\delta}}) g^* + \delta^{t^{\delta}} u$. Any deviation by player i during period t triggers a punishment phase in which player i is initially minimaxed using the strategy from Lemma 4, and then the players continue with a strategy profile with payoffs $v^i(\hat{a})$ that depend on the realized actions during the minmaxing. The continuation payoffs $v^i(\hat{a})$ are chosen so that all players are indifferent among all actions during the minmaxing phase and the overall payoff from the punishment of player i phase is equal to $v^{i,t^{\delta}-t} = (1 - \delta^{t^{\delta}-t}) g^* + \delta^{t^{\delta}-t} u^{i*}$. We choose u and u^{i*} so that they are sufficiently close to u^* and such that for sufficiently high $\delta < 1$, there exists continuation (n-1)-revealing equilibria $\sigma^{u,\delta}$ and $\sigma^{u^{i*,\delta}}$ with payoffs, respectively, $v^{\pi,\delta}(\sigma^{u,\delta}) \leq_{\pi} u$ and $v^{\pi,\delta}(\sigma^{u^{i*,\delta}}) \leq_{\pi} u^{i*}$. Moreover, we need to choose u^{i*so} that no player has incentives not to deviate.

Let $k^* = \frac{100}{1-\alpha^*}$ and find $\epsilon > 0$ so that $B(u^*, 2k\epsilon) \subseteq \operatorname{conint} F_{n-1}^B(\pi)$. Using compactness, one can show that for sufficiently high δ , for each $u \in B(u^*, k\epsilon)$, there exists a strategy profile $\sigma^{u,\delta}$ that induces payoff $v^{\pi,\delta}(\sigma) \preceq_{\pi} u$ and such that $\sigma^{u,\delta}$ is a (n-1)-revealing equilibrium of game $\Gamma(\pi, \delta)$. (It might be necessary to use public randomization if $u^* \notin \operatorname{int} F_{n-1}^B(\pi)$.)

For each player *i*, find $u^{i*} \in B(u^*, k\epsilon)$ so that

$$u^{i*}(\theta_i) \le u^*(\theta_i) - \frac{2\epsilon}{1 - \alpha^*} \text{ for each } \theta_i, \text{ and}$$
(B.1)
$$u^{i*}(\theta_j) \ge u^*(\theta_j) \text{ for each type } \theta_j \text{ of player } j \ne i.$$

For each $t \leq t^{\delta}$ and each player i, let $v^{i,t} = (1 - \delta^t) g^* + \delta^t u^{i*}$. Because of (B.1), for sufficiently high δ , and each player $j \neq i$,

$$v^{i,t}\left(\theta_{j}\right) \ge v^{j,t}\left(\theta_{j}\right) + 2\epsilon. \tag{B.2}$$

Find m^{ϵ} and δ^{ϵ} from Lemma 4. Assume that $m \geq m^{\epsilon}$ and the discount factor $\delta \geq \delta^{\epsilon}$ are high enough so that $(1 - \delta^m) M < \epsilon$, and $(1 - \delta^m) \epsilon > 2 (1 - \delta) M$.

Let $\mu_j^{i,t,*} = \mu_j^{i,v^{i,t}-\epsilon,m,\epsilon}$ be the minmax strategies of players $j \neq i$ from Lemma 4. Let $M_i^{t,*}(\hat{a}^i)$ be the associated payoff vector of player *i* playing action sequence $\hat{a}^i = (a_0^i, \dots, a_{m-1}^i)$. For each sequence of actions \hat{a}^i of player *i* and \hat{a}^{-i} of players -i, define $\hat{a} = (\hat{a}^i, \hat{a}^{-i})$ and payoff vector $v^i(a)$ so that for each type θ_i of player *i*,

$$(1 - \delta^m) M_i^{t,*} \left(\hat{a}^i; \theta_i \right) + \delta^m v^i \left(\hat{a}; \theta_i \right) = v^{i,t} \left(\theta_i \right),$$

and for each type θ_j of player $j \neq i$,

$$(1-\delta)\sum_{s=0}^{m-1}\delta^{s}g_{j}\left(a_{s}^{i},a_{s}^{-i};\theta_{j}\right)+\delta^{m}v^{i}\left(\hat{a},\theta_{j}\right)=v^{i,t}\left(\theta_{j}\right).$$

Notice that because $M_i^{t,*}(a^i; \theta_i) \leq v^{i,t}(\theta_i) - \epsilon$ for each type θ_i of player i,

$$v^{i}(\hat{a},\theta_{i}) \geq v^{i,t}(\theta_{i}) + (1-\delta^{m})\epsilon > v^{i,t}(\theta_{i}) + 2M(1-\delta).$$

Moreover, due to (B.2), for each type θ_i of player $j \neq i$,

$$v^{i}(\hat{a},\theta_{j}) \geq v^{i,t}(\theta_{j}) - (1-\delta^{m}) M > v^{i,t}(\theta_{j}) - \epsilon > v^{j,t}(\theta_{j}) + 2M(1-\delta).$$

We are going to construct strategy profile σ . There are two types of regimes:

48

- Normal(v, t) for each $t \leq t^{\delta}$ and v so that (a) there exists $u \in B(u^*, k\epsilon)$ such that $v = (1 \delta^t) g^* + \delta^t u$, and (b) $v(\theta_i) \geq v^{i,t}(\theta_i) + 2M(1 \delta)$ for each player i and type θ_i . Players play action profile a^* for t periods s = 0, 1, ..., t 1. If there is no deviation, players continue with strategy profile $\sigma^{u,\delta}$. Simultaneous deviations of two or more players are ignored. A deviation by single player i in period s initiates regime Punishment(i, t s).
- Punishment(i, t): The regime lasts m periods. Players -i play strategies $\mu_{-i}^{i,t,*}$. Player i randomizes uniformly across all action sequences $(a_0^i, ..., a_{m-1}^i)$. In particular, the strategies of each player do not depend on their types. After m periods, regime $Normal(v^i(\hat{a}), t)$ is initiated, where \hat{a} are the actions played during the regime.

The profile starts in regime $Normal(v^*, t^{\delta})$.

We compute the payoffs and verify the incentives in the above profile. Initially, we make a preliminary (and perhaps incorrect) assumption that the payoffs in the profiles that end phase $Normal((1 - \delta^t) g^* + \delta^t u, t)$ are equal to u (instead of $v^{\pi,\delta} (\sigma^{u,\delta}) \preceq_{\pi} u$). Then the expected payoff in the beginning of regime Normal(v,t) is equal to v and the expected payoff in the beginning of regime Punishment(i,t) is equal to $v^{i,t}$. Any one-shot deviation during the Normal(v,t) period leads to a payoff not higher than $(1 - \delta) M + \delta v^{i,t}$. If $v(\theta_i) \ge v^{i,t}(\theta_i) + 2M(1 - \delta)$, the deviation is not profitable. In each period of the Punishment(i,t) regime, all players are indifferent among all actions. In particular, they do not have one-shot profitable deviations. Thus, the expected payoff from profile σ under the preliminary assumption is equal to v^* .

Because our preliminary assumption is possibly incorrect, the above argument may not correctly reflect the incentives faced by the players. On one hand, the preliminary assumption does not affect the payoffs of the π -positive probability types. Thus, the behavior prescribed by strategy profile σ is the best response for all such types, given that all positive probability types of the other players follow σ . On the other hand, the behavior prescribed by profile σ may not be the best response for the π -zero probability types. We can modify profile σ so that all the zero-probability types choose the best responses, given the assumption that all (the positive probability types of) other players follow σ . (Notice that this modification does not change the

incentives for the positive probability types.) Because the preliminary assumption may artificially increase the continuation payoffs of the π -zero probability type θ_i , the true expected best response payoffs of this type cannot be higher than $v^*(\theta_i)$. Thus, the true expected payoff from profile σ is equal to $v^{**} \leq_{\pi} v^*$.

Finally, because the strategies prescribe the same (possibly mixed) actions for all π -positive probability types of each player, the beliefs do not get updated before (n-1)-revealing profile σ^u is started.

Lemma 7. If $F_{n}^{A}(\pi) \subseteq F_{n-1}^{+}(\pi)$, then $F_{n}^{B}(\pi) \subseteq F_{n-1}^{+}(\pi)$.

Proof. Take any $v \in F_n^B(\pi)$ and find $\epsilon > 0$ and an incentive-compatible lottery $l = (\alpha, u)$ such that $v = v^{\pi, l}$ and $B(u(a), 2\epsilon) \subseteq \operatorname{int} F_n^A(p^{\pi, l}(a))$ for each positive probability action profile a. We can assume w.l.o.g. that all actions have positive probability.

Using the compactness argument (and possibly public randomization), we can show that there exists δ_0 such that for all $\delta \geq \delta_0$, each a, and each $u' \in B(u(a), \epsilon)$, there exists a strategy profile that induces payoff u' and that is a (n-1)-revealing equilibrium of game $\Gamma(p^{\pi,l}(a,), \delta)$.

For each action profile, let $u^{\delta}(a) = \frac{1}{\delta}u(a) - (1-\delta)g(a) \in B(u(a), \epsilon)$. For each a, find *n*-revealing equilibrium profile σ^a that induces $u^{\delta}(a)$.

Let σ be a strategy profile in which in the first period, players play according to α , and continue with $\sigma(a)$ after first period history a. Then σ is a (n-1)-revealing equilibrium for sufficiently high δ with expected payoff v.

Appendix C. Proof of Theorem 4

In this appendix, we assume that the game has the structure described in section 6.1. In particular,

$$\operatorname{int} IR = \left\{ v \in R^{\Theta^*} : v_i(\theta_i) > 0 \text{ for each type } \theta_i \right\},\$$

and there exist sets $M_i \subseteq R^{\Theta_i}$ such that $\mathbf{0}_i \in M_i$ and the set

int
$$V$$
 = intcon $\left\{ \bigcup_{i} M_{i} \times \{\mathbf{0}_{-\mathbf{i}}\} \right\}$

is not empty. These assumptions imply that the oligopoly games have an open thread.

We begin with a convenient characterization of set $F_1^B(\pi)$.

Lemma 8. Let $v \in R^{\Theta^*}$ be a payoff vector. Then $v \in F_1^B(\pi)$ if and only if for each player *i*, there exist mappings $\beta_i : \Theta \to [0, 1]$ and $m_i : \Theta_i \to M_i$ such that $\sum_i \beta_i^{\theta} \leq 1$ and the following conditions hold:

- (1) Individual rationality: $v_i(\theta_i) > 0$ for each player *i* and type θ_i , and $m_i^{\theta_i}(\theta_i) > 0$ for each player *i* and π -positive probability type θ_i ,
- (2) Incentive compatibility: for all θ_i, θ'_i ,

$$v\left(\theta_{i}\right) \geq m_{i}^{\theta_{i}^{\prime}}\left(\theta_{i}\right) \sum_{\theta_{-i}} \pi^{\theta_{i}}\left(\theta_{-i}\right) \beta_{i}^{\left(\theta_{i}^{\prime},\theta_{-i}\right)},$$

with the equality if type θ_i has π -positive probability and $\theta'_i = \theta_i$.

In particular, set $F_1^B(\pi)$ is convex.

Proof. If v satisfies the above two conditions, then one can easily construct an appropriate lottery to show that $v \in F_1^B$. We show the other direction. Take some $v \in F_1^B(\pi)$ and find π -incentive-compatible lottery $l^0 = (\alpha^0, u^0)$ with value v and such that for each action profile a, either beliefs p(a) are degenerate on the type tuple θ and

$$u^{0}(a) \in F_{1}^{A}\left(\pi^{\theta}\right) = F_{0}\left(\pi^{\theta}\right),$$

or the beliefs p(a) are non-degenerate, and

$$u^{0}(a) \in F_{1}^{A}(p(a)) = \operatorname{int} IR \cap \operatorname{int} V.$$

Because $\operatorname{int} IR \cap \operatorname{int} V \subseteq F_0^B(\pi^\theta)$, we can assume that $u^0(a) \in F_0^B(\pi^\theta)$ for each a played with positive probability by types θ in strategy profile α^0 .

For each π -positive probability type profile θ and action profile a played by positive probability by types in θ , we can find $u^1(a) \in \text{int}V$ such that $u^1(a) \preceq_{\pi^{\theta}} u^0(a)$. Because payoffs $u^0(a)$ are strictly individually rational, we have

$$\max\left\{0, u^{1}\left(\theta_{i}'|a\right)\right\} \leq u^{0}\left(\theta_{i}'|a\right) \text{ for each type } \theta_{i}'.$$

Define allocation $u : \Theta \to \text{int}V$ so that for each type profile θ (not necessarily positive probability),

$$u^{\theta} = \sum_{a} \left(\prod_{i} \alpha_{i}^{0} \left(a_{i} | \theta_{i} \right) \right) u^{1} \left(a \right).$$

For each type profile θ , player i, find $\beta_i^{\theta} \ge 0$ and $\hat{m}_i^{\theta} \in M_i$ so that $\sum_i \beta_i^{\theta} \le 1$ and $u_i^{\theta} = \beta_i^{\theta} \hat{m}_i^{\theta}$. Finally, for each type θ_i , define

$$m_i^{\theta_i} = \frac{\sum_{\theta_{-i}} \pi^{\theta_i} \left(\theta_{-i}\right) u_i^{\left(\theta_i,\theta_{-i}\right)}}{\sum_{\theta_{-i}} \pi^{\theta_i} \left(\theta_{-i}\right) \beta_i^{\left(\theta_i,\theta_{-i}\right)}} = \frac{\sum_{\theta_{-i}} \pi^{\theta_i} \left(\theta_{-i}\right) \beta_i^{\left(\theta_i,\theta_{-i}\right)} \hat{m}_i^{\left(\theta_i,\theta_{-i}\right)}}{\sum_{\theta_{-i}} \pi^{\theta_i} \left(\theta_{-i}\right) \beta_i^{\left(\theta_i,\theta_{-i}\right)}}.$$

Notice that $m_i^{\theta_i}$ is a convex combination of elements of M_i ; hence $m_i^{\theta_i} \in M_i$.

We check that β and m satisfy the thesis of the Lemma. For each π -positive probability type profile $\theta = (\theta_i, \theta_{-i})$, each player i and each action profile $a = (a_i, a_{-i})$ such that a_i is played with positive probability by type θ_i , $u^1(\theta_i|a) = u^0(\theta_i|a) > 0$. It follows that $m_i^{\theta_i}(\theta_i) > 0$ is a convex combination of strictly positive values.

Further, because lottery l^0 is π -incentive compatible, for each action a_i ,

$$v\left(\theta_{i}\right) \geq \sum_{\theta_{-i}} \pi^{\theta_{i}}\left(\theta_{-i}\right) u_{i}^{0}\left(\theta_{i}|a_{i},\alpha_{-i}^{0}\left(\theta_{-i}\right)\right)$$

with equality when action a_i is played with positive probability by type θ_i , i.e., $\alpha_i^0(a_i|\theta_i) > 0$. It follows that for π -positive probability type θ_i

$$v(\theta_{i}) = \sum_{\theta_{-i}} \pi^{\theta_{i}}(\theta_{-i}) u^{0}(\theta_{i}|\alpha_{i}(\theta_{i}), \alpha_{-i}(\theta_{-i}))$$

$$= \sum_{\theta_{-i}} \pi^{\theta_{i}}(\theta_{-i}) u^{1}(\theta_{i}|\alpha_{i}^{0}(\theta_{i}), \alpha_{-i}^{0}(\theta_{-i}))$$

$$= \sum_{\theta_{-i}} \pi^{\theta_{i}}(\theta_{-i}) u^{\theta}(\theta_{i}|\alpha_{i}^{0}(\theta_{i}), \alpha_{-i}^{0}(\theta_{-i}))$$

$$= \sum_{\theta_{-i}} \pi^{\theta_{i}}(\theta_{-i}) \beta^{(\theta_{i},\theta_{-i})} \hat{m}_{i}^{(\theta_{i},\theta_{-i})}(\theta_{i})$$

$$= \sum_{\theta_{-i}} \pi^{\theta_{i}}(\theta_{-i}) \beta^{(\theta_{i},\theta_{-i})} m_{i}^{\theta_{i}}(\theta_{i})$$

and for all types θ_i, θ'_i ,

$$v\left(\theta_{i}\right) \geq \sum_{\theta_{-i}} \pi^{\theta_{i}}\left(\theta_{-i}\right) E_{\alpha^{0}\left(\theta_{i}',\theta_{-i}\right)} u_{i}^{0}\left(\theta_{i}|a\right)$$

$$\geq \sum_{\theta_{-i}} \pi^{\theta_{i}}\left(\theta_{-i}\right) E_{\alpha^{0}\left(\theta_{i}',\theta_{-i}\right)} u_{i}^{1}\left(\theta_{i}|a\right)$$

$$\geq \sum \pi^{\theta_{i}}\left(\theta_{-i}\right) E_{\alpha^{0}\left(\theta_{i}',\theta_{-i}\right)} u_{i}^{1}\left(\theta_{i}|\alpha_{i}^{0}\left(\theta_{i}'\right),\alpha_{-i}^{0}\left(\theta_{-i}\right)\right)$$

$$\geq m_{i}^{\theta_{i}'}\left(\theta_{i}\right) \left(\sum_{\theta_{-i}} \pi^{\theta_{i}}\left(\theta_{-i}\right) \beta_{i}^{\left(\theta_{i}',\theta_{-i}\right)}\right).$$

The last claim follows from the characterization.

52

н		
L		
н		
н		

Take any individually rational vector v^* of payoffs that are individually rational for all positive π -probability types of all players and that can be obtained by a play of nonrevealing actions followed by a payoff vector from stage 1B, $v^* = \gamma g + (1 - \gamma) v'$ for some $g \in V$ and $v' \in F_1^B(\pi)$. The next Lemma shows that there exists a corresponding fully revealing payoff v, with the same payoffs as v^* for the positive probability types and not smaller, and individually rational payoffs for the zero-probability types. The idea is to delay the play of non-revealing actions after the revelation. We need to be careful so that the expected payoffs and the incentives to reveal information truthfully are not affected and that the continuation payoffs after the revelation are individually rational.

Lemma 9. For each $\pi \in \Pi$, $F_2^A(\pi) = F_1^B(\pi)$.

Proof. Take $v^* \in \operatorname{int} \left(IR \cap \operatorname{con} \left(F_1^B(\pi) \cup V \right) \right)$. Find $\gamma_i \geq 0$ and $m_i^* \in M_i$, and $u^* \in F_1^B(\pi)$ so that $\sum_i \gamma_i \leq 1$ and for each player i,

$$v_{i}^{*}\left(\theta_{i}\right) = \gamma_{i}m_{i}^{*}\left(\theta_{i}\right) + \left(1 - \sum_{i}\gamma_{i}\right)u^{*}\left(\theta_{i}\right),$$

with equality for π -positive probability types θ_i . Using Lemma 8, find $\beta_i^{\theta} \ge 0$ and $m_i^{\theta_i} \in M_i$ for each player type tuple θ so that $\sum_i \beta_i^{\theta} \le 1$ and $m_i^{\theta_i} \ge 0$ for each $\theta = (\theta_i, \theta_{-i})$, and

$$v_{i}^{*}(\theta_{i}) = \gamma_{i}m_{i}^{*}(\theta_{i}) + \left(\sum_{\theta_{-i}} \pi^{\theta_{i}}(\theta_{-i})\beta_{i}^{\left(\theta_{i}^{\prime},\theta_{-i}\right)}\right)m_{i}^{\theta_{i}^{\prime}}(\theta_{i})$$
$$\geq \gamma_{i}m_{i}^{*}(\theta_{i}) + \left(\sum_{\theta_{-i}} \pi^{\theta_{i}}(\theta_{-i})\beta_{i}^{\left(\theta_{i}^{\prime},\theta_{-i}\right)}\right)m_{i}^{\theta_{i}^{\prime}}(\theta_{i})$$

with the equality if type θ_i has π -positive probability and $\theta'_i = \theta_i$.

For each player *i* and type profile $\theta = (\theta_i, \theta_{-i})$, let

$$\hat{\beta}_i^{\theta} = \gamma_i + \left(1 - \sum_i \gamma_i\right) \beta_i^{\theta}.$$

For all π -positive probability types θ_i , let

$$\hat{m}_{i}^{\theta_{i}} = \frac{\gamma_{i}m_{i}^{*} + (1 - \sum_{i}\gamma_{i})\sum_{\theta_{-i}^{'}}\pi^{\theta_{i}}(\theta_{-i})\beta_{i}^{(\theta_{i},\theta_{-i}^{'})}m_{i}^{\theta_{i}}}{\gamma_{i} + (1 - \sum_{i}\gamma_{i})\sum_{\theta_{-i}^{'}}\pi^{\theta_{i}}(\theta_{-i})\beta_{i}^{(\theta_{i},\theta_{-i}^{'})}}.$$

For all π -zero probability types θ_i , let $\hat{m}_i^{\theta_i} = \mathbf{0}_i$. For each player *i* type θ_i , define

$$v(\theta_i) = \sum_{\theta_{-i}} \pi^{\theta_i}(\theta_{-i}) \hat{\beta}_i^{(\theta_i,\theta_{-i})} \hat{m}_i^{\theta_i}(\theta_i) \text{ for } \pi\text{-positive prob.}$$
$$v(\theta_i) = v_i^*(\theta_i) \text{ for } \pi\text{-zero prob. } \theta_i.$$

 θ_i

Simple calculations show that $v = v^*$.

We check that assignments $\hat{\beta}_i$ and \hat{m}_i^{θ} satisfy the conditions of Lemma 8 for v. Indeed, $v_i(\theta_i) > 0$ and $\hat{m}_i^{\theta_i} \in M_i$ because $\hat{m}_i^{\theta_i} = \mathbf{0}_i$ or $\hat{m}_i^{\theta_i}$ is a convex combination of elements of M_i . Moreover, for each tuple θ ,

$$\sum_{i} \left(\gamma_{i} + \left(1 - \sum_{i} \gamma_{i} \right) \beta_{i}^{\theta} \right) = \sum_{i} \gamma_{i} + \left(1 - \sum_{i} \gamma_{i} \right) \sum_{i} \beta_{i}^{\theta}$$
$$\leq \sum_{i} \gamma_{i} + \left(1 - \sum_{i} \gamma_{i} \right) \leq 1.$$

The individual rationality holds because, in the first case, $\hat{m}_i^{\theta_i}(\theta_i)$ is equal to $v^*(\theta_i)$ multiplied by a positive factor, and in the second case, $\hat{m}_i^{\theta}(\theta_i) = 0$.

We check the incentive compatibility: for all types π -positive probability type θ_i and all types θ'_i ,

$$\begin{split} v\left(\theta_{i}\right) &= \sum_{\theta_{-i}} \pi^{\theta_{i}}\left(\theta_{-i}\right) \hat{\beta}_{i}^{\left(\theta_{i},\theta_{-i}\right)} \hat{m}_{i}^{\theta_{i}}\left(\theta_{i}\right) \\ &= \gamma_{i} m_{i}^{*} + \left(1 - \sum_{i} \gamma_{i}\right) \sum_{\theta_{-i}^{\prime}} \pi^{\theta_{i}}\left(\theta_{-i}\right) \beta_{i}^{\left(\theta_{i},\theta_{-i}^{\prime}\right)} m_{i}^{\theta_{i}}\left(\theta_{i}\right) \\ &\geq \gamma_{i} m_{i}^{*} + \left(1 - \sum_{i} \gamma_{i}\right) \left(\sum_{\theta_{-i}^{\prime}} \pi^{\theta_{i}}\left(\theta_{-i}\right) \beta_{i}^{\left(\theta_{i}^{\prime},\theta_{-i}^{\prime}\right)}\right) m_{i}^{\theta_{i}^{\prime}}\left(\theta_{i}\right) \\ &= \left(\sum_{\theta_{-i}} \pi^{\theta_{i}}\left(\theta_{-i}\right) \hat{\beta}_{i}^{\left(\theta_{i}^{\prime},\theta_{-i}\right)}\right) \hat{m}_{i}^{\theta_{i}^{\prime}}\left(\theta_{i}\right). \end{split}$$

where the first inequality follows from the choice of β . The incentive compatibility in case of π -zero probability types θ_i is trivial. It follows that $v \in F_1^B(\pi)$.

We can conclude the proof of Theorem 4. The proof is an application of the characterization of the set of equilibrium payoffs from Theorem 3. It is enough to show that $F_2^A(\pi) = F_1^B(\pi)$ and $F_2^B(\pi) = F_1^B(\pi)$. The first claim follows from Lemma 9. The second claim follows from the first and the fact that the composition

of an incentive-compatible lottery with an incentive-compatible and fully revealing lottery can be replaced by a single incentive-compatible lottery with the same value and outcomes that are convex combinations of the outcomes in the original lotteries. The characterization of equilibrium payoffs comes from Lemma 8.

Appendix D. Proof of Lemma 2

On the contrary, suppose that $u^*(\pi)$ is the thread. Let $u^{ns} = u^*\left(\pi^{(\text{normal}_1, \text{ strong}_2)}\right)$ be the thread Nash equilibrium payoff vector, given that the first player is revealed to be normal and the second player is revealed to be strong. Because the equilibrium payoffs must be individually rational, it must be that

 u_1^{ns} (normal₁) ≥ 0 and u_2^{ns} (strong₂) $\geq 1 - 3x$.

By Theorem 5, there exists $\alpha \in \Delta A$ such that

$$u_1^{ns} (\text{normal}_1) = 2\alpha_{WW} + 4\alpha_{TW} - 2\alpha_{TT} \ge 0,$$
(D.1)
$$u_2^{ns} (\text{strong}_2) = 2x\alpha_{WW} + (1+3x)\alpha_{WT} + (1-3x)\alpha_{TT} \ge 1-3x,$$

and

$$u_2^{ns}$$
 (normal₂) $\geq 2\alpha_{WW} + 4\alpha_{WT} - 2\alpha_{TT}$.

The next result shows that u_2^{ns} (normal₂) > 2.

Lemma 10. Suppose that $x \leq \frac{3}{100}$. Then $2\alpha_{WW} + 4\alpha_{WT} - 2\alpha_{TT} > 2$ for each $\alpha \in \Delta A$ that satisfies inequalities (D.1).

Proof. The first inequality in (D.1) implies that

$$\alpha_{TT} \leq \frac{2}{3} - \frac{1}{3}\alpha_{WW} - \frac{2}{3}\alpha_{WT}.$$

Substituting into the second inequality, we obtain

$$2x\alpha_{WW} + (1+3x)\,\alpha_{WT} \ge (1-3x)\left(\frac{1}{3} + \frac{1}{3}\alpha_{WW} + \frac{2}{3}\alpha_{WT}\right),\,$$

or, after some algebra,

$$\alpha_{WT} \ge \frac{1 - 3x}{1 + 15x} + \frac{1 - 9x}{1 + 15x} \alpha_{WW}$$

It follows that

$$2\alpha_{WW} + 4\alpha_{WT} - 2\alpha_{TT}$$

$$\geq \frac{8}{3}\alpha_{WW} + \frac{16}{3}\alpha_{WT} - \frac{4}{3}$$

$$\geq \left(\frac{8}{3} + \frac{16}{3}\frac{1-9x}{1+15x}\right)\alpha_{WW} + \frac{16}{3}\frac{1-3x}{1+15x} - \frac{4}{3} > 2,$$

where the last inequality holds for all $\alpha_{WW} \ge 0$ and all $x < \frac{3}{100}$.

A symmetric argument shows that $u_1^{sn} (\text{normal}_1) > 2$, where u^{sn} is the thread equilibrium payoff vector if the first player is strong, and the second player is normal. Because players must be *ex post* indifferent about revealing their type truthfully, we have

$$u_2^{nn} (\operatorname{normal}_2) = u_2^{ns} (\operatorname{normal}_2) > 2,$$

$$u_1^{nn} (\operatorname{normal}_1) = u_1^{sn} (\operatorname{normal}_1) > 2,$$

where u^{nn} is the thread payoff vector if both players are revealed to be normal.

On the other hand, the sum of the payoffs of the normal types, given any action profile, is never higher than 4. This implies that for any equilibrium payoff vector $u \in NE$ (normal₁, normal₂), u_1 (normal₁)+ u_2 (normal₂) ≤ 4 . The contradiction shows that u^* cannot be a thread.

UNIVERSITY OF TORONTO E-mail address: mpeski@gmail.com.

56