

Efficient chip strategies in repeated games

Wojciech Olszewski and Mikhail Safronov*

February 10, 2018

Abstract

We study a class of chip strategies in repeated games of incomplete information. This class generalizes the strategies studied by Möbius (2001) in the context of a favor-exchange model and the strategies studied in our companion paper Olszewski and Safronov (2017).

In two-player games, if players have private values and their types evolve according to independent Markov chains, then under very mild conditions on the stage game, the efficient outcome can be approximated by chip-strategy equilibria when the discount factor tends to 1.

We extend this result (assuming stronger conditions) to stage games with any number of players. Chip strategies can be viewed as a positive model of repeated interactions, and the insights from our analysis seem applicable in similar contexts, not covered by the present analysis.

1 Introduction

Models of repeated games of incomplete information have a wide range of applications. They include: (a) oligopoly markets in which firms privately know their costs, (b) repeated auctions in which bidders privately know their valuations, and (c) favor exchange when a person in need does not know if others can help.

In the existing literature, repeated games of incomplete information have been analyzed by means of two kinds of strategies: (a) simple and intuitive strategies that allow attainment of only limited payoffs, or (b) strategies that allow attainment of a wider range of payoffs but are less intuitive and more involved, or have been “tailored” with the objective of attaining particular payoffs.

In our companion paper, Olszewski and Safronov (2017), we study a class of simple strategies, called *chip strategies*, and show that equilibria in strategies from this class approximate efficient outcomes in many economic applications, including several models studied extensively in the existing literature, such as the favor-exchange model, repeated auctions, or repeated duopolies. Intuitively, according to these strategies a player takes an individually suboptimal action if that action creates a “gain” for the opponent which is

*Olszewski: Department of Economics, Northwestern University, wo@northwestern.edu; Safronov: Faculty of Economics, University of Cambridge, ms2329@cam.ac.uk

larger than the player’s “loss” from taking it. In exchange, the player implicitly obtains from the opponent a chip (or a fraction of a chip) that entitles the player to receive this kind of favor at some future date. Players are initially endowed with a certain number of chips, and a player who runs out of chips is not entitled to receive any additional favors—and is on suspension—until she provides a favor to her opponent, in which case the player receives one chip back. Detectable deviations result in a breakdown of cooperation.

Perhaps unsurprisingly, Olszewski and Safronov (2017) show that the scope of their simple chip equilibria is limited. Such equilibria cannot approximate efficient outcomes in some (even symmetric) two-player games in which players have more than two types. The aim of this paper is to explore a wider class of chip strategies, which are still relatively simple and intuitive, resemble what we observe in numerous settings in practice, and support the efficient outcomes in a wider range of games.

Although these more general chip strategies are defined for games with any number of players, the idea is probably most intuitive in two-player games. According to these more general chip strategies, a player who takes an action which “helps” the other player increases her chance of implicitly obtaining a chip from the opponent. If, on the other hand, a player takes an action which “hurts” the other player, she increases her chance of giving the opponent a chip.

The current paper differs from Olszewski and Safronov (2017) in that players no longer exchange favors for chips (or fractions thereof) at some constant rate, instead, the probabilities with which players exchange favors and chips are carefully selected to imitate the d’Aspremont and Gerard-Varet (1979) and Arrow (1979) mechanism by using continuation payoffs as transfers. That is, players internalize the effect of their actions on the current payoffs of their opponents by the effect that their actions have on the number of chips they own, and hence on their own continuation payoffs.

For two-player games, we show that if players’ types are i.i.d. or more generally evolve according to independent Markov chains, then under some very mild conditions on the stage game, the efficient outcome can be approximated by chip-strategy equilibria when the discount factor tends to 1. The fact that the efficient outcome can be approximated is roughly intuitive. When players assign higher weights to future payoffs, one can equip them with a larger number of chips. Thus, players go on inefficient suspension less frequently, and the action profile that maximizes the aggregate payoff is played more frequently.

With more than two players, additional complexity arises if a favor is provided by more than one player and more than one player benefits from this favor. We provide efficiency results for this case as well, but under somewhat stronger assumptions. This suggests that chip strategies are more natural for bilateral interactions, although chip strategies can still be used for multilateral interactions in a wide range of applications.

Related literature

Fudenberg, Levine, and Maskin (1994) prove a folk theorem for a family of repeated games in which players have i.i.d. types and private values. The focus of their paper is entirely on the payoffs that can be attained in equilibria, not on the strategies that yield these payoffs. Escobar and Toikka (2013) study the more general model in which players have Markov independent types and private values. They show that

the efficient payoffs could be attained in some version of review strategies; moreover, Escobar and Toikka prove that any Pareto-efficient payoff vector above a stationary minmax vector can be attained for a generic class of games. The review strategies¹ used by Escobar and Toikka are intuitive and deliver general results. However, their complete strategies are not entirely explicit. Contingent on some histories, they are defined by a fixed-point argument. Hörner, Takahashi, and Vieille (2015) provide even more general results, relaxing the assumptions of private values and independent types; the focus of their paper is also on the payoffs that can be attained in equilibria.²

Some versions of chip strategies were introduced in the context of a two-player, favor-exchange model by Möbius (2001) and studied in the subsequent papers by Hauser and Hopenhayn (2008) and Abdulkadiroglu and Bagwell (2012). Some form of one-chip strategies also appears in the study of oligopolies (see Athey, Bagwell, and Sanchirico (2004), Athey and Bagwell (2008), and the discussion of repeated games with incomplete information in Mailath and Samuelson (2006)), and in the papers on repeated auctions by Aoyagi (2003, 2007), and Rachmilevitch (2013). Chip strategies have also been used by Wolitzky (2015) in a model of information transmission on networks.

The focus of our paper is on behavior rather than on payoffs. We do not pretend to improve on the existing literature in terms of generality of efficiency results. Indeed, Fudenberg, Levine, and Maskin, and Escobar and Toikka provide quite general results for stage games with private values and independent types. Hörner, Takahashi, and Vieille (2015) provide even more general results, relaxing the assumptions of private values, independent types, and perfect monitoring. We view our main contribution as describing intuitive strategies which provide a positive model of playing repeated games with incomplete information. Under relatively mild conditions, these strategies approximate efficient payoffs in settings with Markov independent types and private values. However, our methods also allow for showing the existence of efficient equilibria in settings in which this existence cannot be derived from the literature, namely, in some repeated auctions with correlated types in which the monitoring is imperfect, and the values may not be private. (Moreover, the construction of efficient equilibria is much simpler for the repeated auctions. And so we relegate the analysis to Olszewski and Safronov (2017), in which we study such simpler strategies.)

The rest of the paper is organized as follows. In Section 2, we introduce the model and present our main assumptions on the stage game. In Section 3, we show the efficiency result for repeated games with transfers. This result is not new; the purpose of including it is to demonstrate the idea of transfers as in the d'Aspremont and Gerard-Varet (1979) and Arrow (1979) mechanism, which will later be imitated by continuation payoffs. In Sections 4 and 5, we describe chip-strategy equilibria for two-player games with i.i.d.

¹Review strategies were initially studied by Radner (1985) in a repeated moral-hazard game. In the case of i.i.d. costs, the ideas behind Escobar and Toikka's equilibria are also closely related to the linking mechanism from Jackson and Sonnenschein (2007).

²They also imitate the d'Aspremont and Gerard-Varet (1979) and Arrow (1979) mechanism to construct suitable monetary transfers which are recursively imitated by continuation payoffs.

types and prove their approximate efficiency. We generalize this result to two-player games with Markov types and games with any number of players in Sections 6 and 7, respectively. We discuss some advantages and drawbacks of chip-strategy equilibria in Section 8.

2 Preliminaries

2.1 Model

Consider a normal-form game G with I players, indexed by $i = 1, \dots, I$. Let A_i and Θ_i be the finite sets of actions and (privately observed) types, respectively, of player i . Let $u_i(\theta_i, a)$ be the payoff of player i . We make some assumptions on the payoffs, but it will be convenient to postpone presenting them to the end of this section.

We study a repeated game in which players play stage game G in periods $t = 1, 2, \dots$, and discount future payoffs at a common rate δ ; it is convenient to denote $1 - \delta$ by ε . Actions are publicly observed at the end of each period. In the repeated game, players are allowed to communicate at the beginning of each period by sending simultaneous, publicly observed cheap-talk messages regarding their types. We assume that the message space of each player i coincides with the type space Θ_i .³ Players also have access to a public randomization device, that is, they observe the realization of a random variable distributed uniformly on the interval $[0, 1]$. The timing of events in each period is as follows: (a) players privately observe their types; (b) they simultaneously send public cheap-talk messages regarding their types; (c) players simultaneously take publicly observed actions; (d) they observe a realization of the public randomization device.

We will first assume that players' types are i.i.d. according to distributions η_i , $i = 1, \dots, I$, and then generalize the results to the case in which players' types are still independently distributed but evolve over time according to homogeneous, aperiodic irreducible Markov chains. We assume that there exists a t such that for every pair of type profiles θ, θ' , if the type profile in the current period is θ , the type profile t periods from now will be θ' with positive probability. By the Ergodic Theorem (see, for example, Shiryaev, 1996), every such process has a limiting type distribution $\eta = (\eta_1, \dots, \eta_I)$, and independently of the initial type profile, the distribution of types at time t converges to the limiting distribution at an exponential rate as $t \rightarrow \infty$.

All other elements of the model, that is, histories, repeated-game strategies, and payoffs, are defined in the standard manner.

2.2 Assumptions on the stage game

In this section, we present the assumptions that we impose on the stage game in the case of two players with i.i.d. types. Since these assumptions will also be used in the case of more than two players, we present them

³This message space can be replaced with any other message space that has at least as many number of elements.

in the general form. We will adjust these assumptions to the Markov case in Section 6. For more than two players, we will need an additional assumption, which we introduce and discuss in Section 7.

For any set of players $R \subset \{1, \dots, I\}$ and their type profile $\theta_R \in \Theta_R = \prod_{i \in R} \Theta_i$, denote by $a(\theta_R)$ an *R-efficient action profile*, that is, an action profile that maximizes the total payoff of all players in R . For any θ , denote by $a(\theta)$ the *efficient action profile*, that is, the action profile that maximizes the sum of the stage-game payoffs of all I players. We will make the following assumptions:

Assumption I: For all subsets $R \subset \{1, \dots, I\}$ and all type profiles $\theta_R \in \Theta_R$ there is a unique action profile $a(\theta_R)$ that maximizes the total payoff of all players in R , given the profile θ ; in addition, $a(\theta_R) \neq a(\theta'_R)$ for all θ_R and θ'_R such that $\theta_i \neq \theta'_i$ for some $i \in R$ and $\theta_j = \theta'_j$ for all $i \neq j \in R$.

For $i \notin R$ let $v_R^i = E_\theta(u_i(\theta_i, a(\theta_R)))$ denote the expected stage-game payoff of player i , when players take the *R-efficient action profile*. Similarly, for $i \in R$ let w_R^i denote the expected stage-game payoff of player i , when players take the *R-efficient action profile*. We will sometimes call the players in R *active* and say that the players not in R are *on suspension*.

Assumption II: For any $i, j \in \{1, \dots, I\}$ and any R such that $i, j \in R$,

$$v_{R-\{i\}}^i < w_R^i < w_{R-\{j\}}^i.$$

Here and throughout the paper, when we use symbols i and j to denote players, we assume that $i \neq j$, that is, that these two letters denote different players. Assumption II guarantees that every player prefers being active to being on suspension (regardless of which of the other players are active), and that every player gains from another player going on suspension.

Assumption III (in the i.i.d. case): The incomplete information stage game has a Bayesian Nash equilibrium in which the payoff of every player i is strictly less than w_R^i for $R = \{1, \dots, I\}$.

We will call the equilibrium described in Assumption III the *bad equilibrium*. Though Assumption III does not hold for all games, bad equilibria do exist in many settings of interest. In addition, in symmetric games a stage-game symmetric equilibrium always exists. This equilibrium can be either efficient, or inefficient. In the latter case, the game satisfies Assumption III and allows for the construction of an efficient equilibrium in chip strategies.

3 The main idea

We first focus on the i.i.d. types. Later, we generalize our results to the case when players' types are Markovian. The main idea of our construction is to imitate the AGV mechanism (see d'Aspremont and

Gerard-Varet (1979) and Arrow (1979)) using variations in continuation payoffs as transfers. To introduce this idea, we assume in this section that players are allowed to make monetary transfers to one another at the end of each period, and that these transfers enter players' payoffs in the quasi-linear manner. We construct an efficient equilibrium using monetary transfers. Then, in the next section, we will show that those monetary transfers can be imitated by variations in continuation payoffs.

Theorem 1. *If the stage game satisfies Assumption III and players are allowed to make monetary transfers at the end of each period, then the efficient payoff can be attained in equilibrium if the discount factor is sufficiently large. If in addition the game satisfies Assumption I, then the efficient payoff can be attained in strict equilibrium.*

We prove Theorem 1 by essentially constructing the expected externality mechanism of d'Aspremont and Gerard-Varet. The claim in Theorem 1 is not new; for example, Athey and Segal (2013) established a more general result (see their Proposition 2). Our main objective, however, is to prove the result without transfers, and in chip strategies. We provide the proof of Theorem 1 not only for completeness, but we will refer to this construction in the following sections.

First, we introduce some auxiliary terms, which will also be used later. Let

$$s_i^j = E_{\theta_{-i}}(u_j(\theta_j, a(\theta_i, \theta_{-i}))) - E_{\theta_i} E_{\theta_{-i}}(u_j(\theta_j, a(\theta_i, \theta_{-i})))$$

be the expected effect of player i 's report θ_i on player j 's payoff; in particular, $s_i^j > 0$ ($s_i^j \leq 0$) if player i reports a type that gives player j in expectation a payoff higher (no higher) than the ex ante expected payoff. This effect is obviously a function of θ_i , but we will often disregard its argument as that will cause no confusion. Let

$$s_i = \sum_{j \neq i} s_i^j$$

be the effect of player i 's report on the total payoff of all other players. Further, let

$$p_i = \Pr\{s_i > 0\} \cdot E_{\theta_i}[s_i | s_i > 0],$$

or equivalently,

$$p_i = -\Pr\{s_i \leq 0\} \cdot E_{\theta_i}[s_i | s_i \leq 0].$$

To show Theorem 1, consider the following strategies.

- (A) In every period, players report their types truthfully.
- (B) If θ is the reported type profile, players take action profile $a(\theta)$.
- (C) Players make transfers. For all $i, j \in \{1, \dots, I\}$, player j transfers s_i^j to player i .

That is, player i obtains (as a transfer) the difference between the sums of the interim and ex ante expected payoffs of the other players. Player i 's expected payoff from reporting θ_i' , given truthful reporting

of the other players, is then

$$E_{\theta_{-i}}(u_i(\theta_i, a(\theta'_i, \theta_{-i}))) + \sum_{j \neq i} [E_{\theta_{-i}}(u_j(\theta_j, a(\theta'_i, \theta_{-i}))) - E_{\theta_i} E_{\theta_{-i}}(u_j(\theta_j, a(\theta_i, \theta_{-i})))] - \sum_{j \neq i} E_{\theta_j} [E_{\theta_{-j}}(u_i(\theta_i, a(\theta_i, \theta_{-i}))) - E_{\theta_j} E_{\theta_{-j}}(u_i(\theta_i, a(\theta_i, \theta_{-i})))].$$

The first term in this expression is player i 's expected interim payoff given her actual and reported types, the second term is the expected payment to player i from the other players, and the third term is the expected payment of player i to the other players. The third term is equal to zero, and the second part of the second term does not depend on player i 's report, while the first term and the first part of the second term sum up to

$$\sum_{j=1}^I E_{\theta_{-i}}(u_j(\theta_j, a(\theta'_i, \theta_{-i}))) \quad (1)$$

Thus, if the players other than i report truthfully, player i has incentives to maximize the sum of the stage-game payoffs, which is attained by reporting her own type truthfully.

(D) An action profile other than $a(\theta)$ for any reported type profile θ , or any refusal to make the prescribed transfers, triggers a permanent repetition of the bad stage-game equilibrium.

This obviously disciplines the players to take action profile $a(\theta)$ given any report θ , and to make the prescribed transfers.

The prescribed strategies are incentive compatible and attain the efficient payoff.

4 Efficient chip strategies for two players with i.i.d. types

We will now specify chip strategies that approximate the efficient outcome in an arbitrary two-player game that satisfies our assumptions: At the beginning of each period, player i holds $k_i \in \{0, \dots, 2n\}$ chips, where $k_1 + k_2 = 2n$, with $k_1 = k_2 = n$ at the beginning of the repeated game. The number k_i is the state of the game. The game is played as follows:

(A) is as in Section 3, that is, players report their types truthfully. (B) is almost as in Section 3, that is, if θ is the reported type profile and $k_i \neq 0, 2n$, then players take action profile $a(\theta)$. We will specify in (E) what happens in states $k_i = 0, 2n$.

The state for the following period is determined at the end of the current period, contingent on the realization of the public randomization device, by the following four-component lottery. The first component describes the way in which player i 's report affects the state. The second component describes the way in which the report of player j affects the state. These two lotteries are independent. The third and fourth components are adjustment terms, independent of the reports. Suppose player i currently holds $k_i = k \neq 0, 2n$ chips. Then (C) from Section 3 is replaced with (C1)–(C4):

(C1) Player i obtains a chip from player j with probability $\alpha_k^i s_i \mathbf{1}_{\{s_i > 0\}}$ and gives player j a chip with probability $-\phi_k^i s_i \mathbf{1}_{\{s_i \leq 0\}}$;⁴

The coefficients α_k^i and ϕ_k^i will be specified later. For now, it is important to know that they will converge to 0 as the discount factor converges to 1, which makes these probabilities well-defined.

(C2) Player j gives player i a chip with probability $-\phi_{2n-k}^j s_j \mathbf{1}_{\{s_j \leq 0\}}$ and obtains a chip from player i with probability $\alpha_{2n-k}^j s_j \mathbf{1}_{\{s_j > 0\}}$.

(C3) Player i gives player j a chip with probability

$$\frac{1}{3} - \phi_k^i p_i - \alpha_{2n-k}^j p_j,$$

independently of the messages sent in the current period, and obtains a chip from player j with probability

$$\frac{1}{3} - \alpha_k^i p_i - \phi_{2n-k}^j p_j,$$

independently of the messages sent in the current period;

(C4) The state in the following period is unaltered with the remaining probability.

That is, in the four-component lottery the total probability of player i giving a chip to player j (as well as receiving a chip from player j) is a sum of the corresponding probabilities in (C1)–(C3). With the remaining probability the state is unaltered.

(D) As in Section 3, if an action profile other than $a(\theta)$ is observed, for any reported type profile θ , in the following period players switch to permanently playing the bad stage-game equilibrium.

(E) If player i is left with no chip, she goes on suspension. This means that for the following M periods (where M is defined later) the players report their types truthfully and play the action profile which is most preferred by i 's opponent. After the M periods, player i comes back from suspension, which means that she obtains one chip from the opponent.^{5 6}

Theorem 2. *In any two-player repeated game in which the players' types are i.i.d. and the stage game satisfies Assumptions I–III, the efficient payoff can be arbitrarily closely approximated by chip-strategy strict equilibria when the discount factor δ is sufficiently close to 1.*

The intuition behind our construction and this result is as follows. At the beginning of each period, the transition probabilities between various states are exogenously specified, independently of the players' reports. (In the particular case of our strategies, the probability of transiting to each “neighbor” state is

⁴The symbol $\mathbf{1}_{\{\cdot\}}$ denotes the characteristic function of the set $\{\cdot\}$. That is, $\alpha_k^i s_i \mathbf{1}_{\{s_i > 0\}} = \alpha_k^i s_i$ when $s_i > 0$, and $\alpha_k^i s_i \mathbf{1}_{\{s_i > 0\}} = 0$ when $s_i \leq 0$; similarly, $-\phi_k^i s_i \mathbf{1}_{\{s_i \leq 0\}} = -\phi_k^i s_i$ when $s_i \leq 0$ and $-\phi_k^i s_i \mathbf{1}_{\{s_i \leq 0\}} = 0$ when $s_i > 0$.

⁵It is worth pointing out that truthful reporting is incentive compatible for players on suspension, since their reports have no effect on the action profiles that are most preferred by their opponents (by the assumption that each player's payoffs are independent of the other players types).

⁶The suspension enlarges the number of states, since at $k = 0$ or $k = 2n$, the strategies depend on the time spent in suspension. For the sake of brevity, we postpone the discussion of additional suspension states to Section 5.2.

1/3, as is the probability of remaining in the current state.) These ex ante probabilities must have the property that the long-run occupation probabilities of the extreme states, in which one of the players goes on suspension, are negligible. This will guarantee the approximate efficiency of our strategies.

These initially specified transition probabilities are then adjusted contingent on players' reports. By imposing a positive externality on the opponent's flow payoff, a player increases her chance of obtaining a chip; and by imposing a negative externality on the opponent's flow payoff, she increases her chance of giving away a chip. Assumption II implies that players prefer having more chips to having fewer chips. More specifically, the adjusted probabilities are specified to imitate the AGV transfers, which guarantees that the players have incentives to maximize the aggregate payoff.

It is possible to find alternative chip strategies which achieve efficiency and induce truth-telling. The adjustment terms in condition (C3) can be changed, making the probability with which a player obtains an additional chip or gives a chip to the opponent different from 1/3. As we go through our analysis, we will indicate two properties of the adjustment terms that are required for our construction, and highlight these two properties in bold font. Similarly, we will indicate and highlight the required properties of α 's and ϕ 's.

5 Analysis

5.1 Value functions

For two-player games, we adopt a slightly simpler notation. Namely, let v^i denote the (expected) stage-game payoff of player i when she is on suspension (i.e., $v^i = v_R^i$ for $R = \{1, 2\} - \{i\}$), and let w_j^i denote the stage-game payoff of player i when player j is on suspension (i.e., $w_j^i = w_R^i$ for $R = \{1, 2\} - \{j\}$), and let w^i denote the stage game payoff of player i when both players are active (i.e., $w^i = w_{\{1,2\}}$).

Denote by V_k^i the continuation payoff of player i in the state in which she has $k_i = k$ chips. These payoffs are computed assuming that players play the prescribed strategies, and at the ex-ante stage when players have not yet learned their current types. We will often call V_k^i 's *value functions*.

The coefficients α_k^i and ϕ_k^i will sometimes be called *probabilities of control*, since they determine the chance that player i 's report will affect the state in the following period. They will be defined in such a way that the following equations are satisfied:

$$\alpha_k^i s_i (1 - \varepsilon) [V_{k+1}^i - V_k^i] = s_i \varepsilon$$

for $s_i > 0$, and

$$\phi_k^i (-s_i) (1 - \varepsilon) [V_{k-1}^i - V_k^i] = s_i \varepsilon$$

for $s_i \leq 0$.⁷

This choice of α_k^i and ϕ_k^i gives player i the incentives to maximize the sum of the stage-game payoffs of both players, in exactly the same way that transfers do in Section 3. Indeed, by (C1)

⁷The value of V_k^i will not depend on the coefficients α_k^i , ϕ_k^i , and will be determined beforehand.

and (C2) of the definition of strategies, player i 's report affects her continuation payoff at the beginning of the following period through the left-hand sides of the equations that define α_k^i and ϕ_k^i . The right-hand sides of those equations are equal to the sum of stage-game payoffs across all players other than i , which, together with the effect of player i 's report on her own stage-game payoff, yields the desired incentives to maximize the sum of the stage-game payoffs of all players.

Of course, α_k^i 's and ϕ_k^i 's may not necessarily imitate the d'Aspremont and Gerard-Varet transfers exactly. **It is sufficient to guarantee that players prefer reporting their types truthfully to reporting other types. Since by Assumption I a unilateral misreport induces an inefficient choice of the action profile, players have strict incentives to report truthfully under the d'Aspremont and Gerard-Varet transfers. It is therefore sufficient to imitate the d'Aspremont and Gerard-Varet transfers only approximately.** In general, however, players may have many type profiles that induce payoffs close to the payoffs induced by truthful reports, and then α_k^i 's and ϕ_k^i 's must be chosen very close to those that exactly imitate the d'Aspremont and Gerard-Varet transfers.

By dividing the equations for α_k^i and ϕ_k^i by s_i , we obtain

$$\alpha_k^i(1 - \varepsilon) [V_{k+1}^i - V_k^i] = \varepsilon \text{ and } \phi_k^i(1 - \varepsilon)[V_k^i - V_{k-1}^i] = \varepsilon. \quad (2)$$

Given the prescribed strategies, value V_k^i for $k = 1, 2, \dots, 2n - 1$ satisfies the following recursive equation:

$$V_k^i = \varepsilon w^i + (1 - \varepsilon)\frac{1}{3}V_{k-1}^i + (1 - \varepsilon)\frac{1}{3}V_k^i + (1 - \varepsilon)\frac{1}{3}V_{k+1}^i. \quad (3)$$

Indeed, player i 's current stage-game payoff is w^i . By (C1) of the definition of strategies, player i gives the opponent a chip with probability $-s_i\phi_k^i$ when $s_i \leq 0$, in expectation, this yields $\phi_k^i p_i$. By (C2) of the definition, player i gives the opponent a chip with probability $s_j\alpha_{2n-k}^j$ when $s_j > 0$, in expectation, this yields $\alpha_{2n-k}^j p_j$. Together with the report-independent chance of giving the opponent a chip described in (C3), this yields a probability of $1/3$ for player i giving a chip, and hence the second term in the expression for V_k^i . In a similar manner, we compute the third term - the chance of staying in the state with $k_i = k$ chips, and the fourth term - the chance of player i obtaining a chip from the opponent.

For $k = 0$ (omitting terms of order smaller than ε),

$$V_0^i = M\varepsilon v^i + (1 - M\varepsilon)V_1^i, \quad (4)$$

and for $k = 2n$ (again omitting terms of order smaller than ε),

$$V_{2n}^i = M\varepsilon w^i + (1 - M\varepsilon)V_{2n-1}^i. \quad (5)$$

5.2 Efficiency

In this section, we will show that the strategies described above approximate the efficient outcome. The strategies induce a stochastic Markov chain over states $k = 2n + M - 1, \dots, 2n, 2n - 1, \dots, 1, 0, \dots, -M + 1$.

States $2n + M - 1, \dots, 2n$ correspond to the periods in which player 2 is on suspension. States $0, \dots, -M + 1$ correspond to the periods in which player 1 is on suspension. States $1, \dots, 2n - 1$ correspond to the periods in which neither player is on suspension, and they indicate the number of chips held by player 1. By the Ergodic Theorem (see, for example, Chapter 1, §12, Theorem 1 in Shiryaev, 1996), there exists a probability distribution π_k over states $\{k = 2n + M - 1, \dots, -M + 1\}$ such that the probability of being in state k after a sufficiently large number of periods is arbitrarily close to π_k , independent of the initial state. This probability distribution π_k is an eigenvector of the transition matrix

$$\begin{bmatrix} 0 & 1 & & 0 & & & & 0 \\ & \cdot & \cdot & & & & & \\ & & \cdot & 1 & & & & \\ & & & & 1/3 & & & 0 \\ 1 & \cdot & \cdot & 0 & \begin{bmatrix} 1/3 & 1/3 & 0 & \cdot & \cdot & \cdot & \cdot \\ 1/3 & 1/3 & 1/3 & \cdot & \cdot & \cdot & \cdot \\ 0 & 1/3 & 1/3 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & 1/3 & 1/3 & 0 \\ \cdot & \cdot & \cdot & \cdot & 1/3 & 1/3 & 1/3 \\ \cdot & \cdot & \cdot & \cdot & 0 & 1/3 & 1/3 \end{bmatrix} & 0 & \cdot & \cdot & 1 \\ & & & 0 & & & & 1/3 & 0 & & & \\ & & & & & & & & & & 1 & \cdot \\ & & & & & & & & & & \cdot & \cdot \\ & & & 0 & & & & & & & & 1 & 0 \end{bmatrix},$$

that corresponds to eigenvalue 1, where the entry in row l and column m is the probability of transiting from state m to state l . It is easy to verify that the eigenvector corresponding to eigenvalue 1 must have

$$\pi_k = \frac{3}{2M + 3(2n - 1)}$$

for $k = 1, \dots, 2n - 1$, and

$$\pi_k = \frac{1}{2M + 3(2n - 1)}$$

for $k = 2n + M - 1, \dots, 2n$ and $k = 0, \dots, -M + 1$.

Since the expected payoff vector in any state $k \in \{1, \dots, 2n - 1\}$ is efficient, it follows that, when δ is sufficiently close to 1, the player's payoff vector is efficient with probability arbitrarily close to

$$\frac{3(2n - 1)}{2M + 3(2n - 1)},$$

which is the sum of the ergodic occupation probabilities of states $k = 1, \dots, 2n - 1$. This probability converges to 1 when n converges to ∞ faster than M does.

Notice that the efficient outcome can be approximated only when the transition probabilities are such that **the ergodic occupation probabilities of states $2n$ and 0 converge to 0**.

5.3 Probabilities of control, incentives

The probabilities of control, α_k^i and ϕ_k^i , are defined by the equations in (2), where the value functions are determined recursively by the equations (3)–(5). Notice that the value functions determined by these equations may not be the actual value functions in the repeated game, since we omitted terms of order smaller than ε in (4) and (5). In particular, the probabilities of control imitate the d’Aspremont and Gerard-Varet transfers only in approximation. Nevertheless, due to Assumption I, players will be induced to report truthfully for sufficiently small values of ε , since they would have strict incentives under the d’Aspremont and Gerard-Varet transfers. Note that the issue of providing incentives concerns only deterring on equilibrium path deviations, that is, that players have incentives to report truthfully. “Off-path deviations” are deterred by the threat of switching to the bad equilibrium.

All that remains to be shown is that the coefficients α_k^i and ϕ_k^i are positive but small. This will be achieved by considering the value of having one more chip,

$$\Delta_k^i := V_{k+1}^i - V_k^i,$$

for $k = 2n-1, 2n-2, \dots, 0$, and proving that all Δ_k^i 's are much larger than ε when M is sufficiently large.

By (3),

$$\Delta_k^i = (1-\varepsilon)\frac{1}{3}\Delta_{k+1}^i + (1-\varepsilon)\frac{1}{3}\Delta_k^i + (1-\varepsilon)\frac{1}{3}\Delta_{k-1}^i$$

for $k = 2n-2, \dots, 1$; and by (4) and (5),

$$\Delta_{2n-1}^i = M\varepsilon w_j^i - M\varepsilon w^i$$

and

$$\Delta_0^i = M\varepsilon w^i - M\varepsilon v^i,$$

where we use the fact that V_{2n-1}^i and V_1^i converge to w^i as ε converges to 0.

For $\varepsilon = 0$, this system of equations is satisfied by all Δ 's that are equal to 0. By the Implicit Function Theorem, Δ 's are differentiable functions of ε . By taking the derivatives of the equations for Δ 's with respect to ε , and plugging in 0 for ε and all Δ 's, we obtain a system of equations for the derivatives of Δ 's at $\varepsilon = 0$. This implies that if we replace each Δ_k^i by its derivative $\partial\Delta_k^i/\partial\varepsilon$ at $\varepsilon = 0$, then our system of equations must be satisfied for $\varepsilon = 0$, and for the free terms $M\varepsilon w_j^i - M\varepsilon w^i$ and $M\varepsilon w^i - M\varepsilon v^i$ replaced with $Mw_j^i - Mw^i$ and $Mw^i - Mv^i$, respectively. In matrix notation, this new system of linear equations can be expressed as

$$\begin{bmatrix} 1 & 0 & 0 & \cdot & \cdot & \cdot & \cdot \\ -1/3 & 2/3 & -1/3 & \cdot & \cdot & \cdot & \cdot \\ 0 & -1/3 & 2/3 & -1/3 & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & -1/3 & 2/3 & -1/3 \\ \cdot & \cdot & \cdot & \cdot & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \partial\Delta_{2n-1}^i/\partial\varepsilon \\ \partial\Delta_{2n-2}^i/\partial\varepsilon \\ \cdot \\ \cdot \\ \cdot \\ \partial\Delta_0^i/\partial\varepsilon \end{bmatrix} = \begin{bmatrix} Mw_j^i - Mw^i \\ 0 \\ \cdot \\ \cdot \\ \cdot \\ 0 \\ Mw^i - Mv^i \end{bmatrix}.$$

By using Gauss-Jordan elimination method, we reduce the system to

$$\begin{bmatrix} 1 & 0 & 0 & \cdot & \cdot & \cdot & \cdot \\ 0 & 2/3 & -1/3 & \cdot & \cdot & \cdot & \cdot \\ 0 & 0 & 1/2 & -1/3 & \cdot & \cdot & \cdot \\ 0 & 0 & 0 & 4/9 & -1/3 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & 0 & (2n-1)/3(2n-2) & -1/3 \\ \cdot & \cdot & \cdot & \cdot & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \partial\Delta_{2n-1}^i/\partial\varepsilon \\ \partial\Delta_{2n-2}^i/\partial\varepsilon \\ \cdot \\ \cdot \\ \cdot \\ \cdot \\ \partial\Delta_0^i/\partial\varepsilon \end{bmatrix} = \begin{bmatrix} M(w_j^i - w^i) \\ M(w_j^i - w^i)/3 \\ M(w_j^i - w^i)/6 \\ M(w_j^i - w^i)/9 \\ \cdot \\ M(w_j^i - w^i)/3(2n-2) \\ M(w^i - v^i) \end{bmatrix}.$$

Let u be the smaller of the two numbers $w_j^i - w^i$ and $w^i - v^i$; both of these numbers are positive, by Assumption II. This implies by recursion (beginning from $k = 0$) that

$$\partial\Delta_k^i/\partial\varepsilon \geq Mu/2$$

for all k . This implies, in turn, that Δ_k^i is larger than $Mu\varepsilon/4$ for all k when ε is sufficiently small.

Notice that players can be provided incentives to maximize the total payoff only when the transition probabilities are such that **the value of having one more chip is high enough compared to ε** .

6 Efficient chip strategies for two players with Markov types

In this section, we generalize the result from the previous section to players' types being independently distributed across players but evolving over time according to homogeneous, aperiodic irreducible Markov chains. We denote the limiting (ergodic) distribution of type profiles by η ; that is, independently of the initial type profile, the distribution of type profiles at time t converges to η (at an exponential rate) as $t \rightarrow \infty$.

To construct the efficient chip strategies for players with Markov types, we will modify the chip strategies constructed in the i.i.d. case, for players' types being distributed according to η . Recall the intuition from the i.i.d. case. The transition probabilities (that is, the probabilities of chip transfers between players) were first defined independently of players' reports, and then adjusted contingent on players' reports to make players internalize the externality they impose on the opponents' flow payoffs.

In the Markov case, players' expected stage-game payoffs depend on the reports from the previous period. Thus, to align the players' individual incentives with the objective of maximizing the total payoff of the two players, the transition probabilities must be adjusted contingent on the reports from the previous period. This causes players' current reports to have an effect on the strategies and payoffs in the following period; and by iterating this argument, in any number of future periods.

However, this effect vanishes at an exponential rate. When δ is close to 1, the transition probabilities can still be adjusted to make players internalize the effect of their reports on other players' payoffs for any given number of future periods. The effects in more remote periods have, however, negligible impact on the current incentives. Notice that strict incentives of each player to report truthfully, guaranteed by Assumption I, are

essential for this argument; if $a(\theta_R) = a(\theta'_R)$ for two profiles θ_R and θ'_R which differ only by the type of player i , then player i could be given incentives to misreport her type even by small differences in her payoff coming from remote periods.

Notice also that the transition probabilities cannot be adjusted to make players internalize the effect of their reports on other players' payoffs in all future periods. The adjustment in the transition probabilities for a finite number of periods is an expression of order $O(\varepsilon)$, while if the effects in all periods were included, then the adjustment would be an expression of order $O(1)$.

In the Markovian case, we retain Assumptions I–III and we use the notations v_R^i, w_R^i for the payoffs as defined for players' types being i.i.d. according to the limiting distribution η . However, we modify Assumption III as follows:

Assumption III (in the Markov case): The repeated game has an equilibrium in which the payoff of every player i is less than w_R^i for $R = \{1, \dots, I\}$.

When types are Markov, we no longer have a repeated-game equilibrium, which would be a repetition of the stage-game strategies. This problem has been pointed out in earlier papers (see, for example, Athey and Bagwell (2008) and Escobar and Toikka (2013)). Thus, we need to assume the existence of a bad equilibrium in the repeated game. It is relatively easy to explicitly construct bad repeated-game equilibria in many concrete settings (such equilibria, called worst carrot-and-stick equilibria, have been constructed for the repeated version of Spulber's oligopoly by Athey and Bagwell (2008)).

6.1 Modified strategies

Assumptions I–III allow for construction of chip-strategy equilibria in the case when players' types are i.i.d. according to η . As before, we use the notation w^i, v^i, w_j^i for the expected stage-game payoffs of player i , when players' types are distributed according to η , and when no player, player i , or player j , respectively, is on suspension. Further, we denote by V_k^i the value functions in this i.i.d. case, and by α_k^i and ϕ_k^i the probabilities of control. We will now adapt these strategies to the Markovian case.

First, we make a more specific assumption about the values of M and n ; we no longer merely require M to tend to infinity more slowly than n ; we also require n/M to tend to infinity more slowly than M . For concreteness, assume that n tends to ∞ at the same rate as $M^{3/2}$:

$$M^{3/2} \sim n. \tag{6}$$

We keep the same probabilities of control, α_k^i and ϕ_k^i , as in the case with i.i.d. types distributed according to η . The only change we make in the construction of strategies is that we define s_i and p_i differently. To define them, we need to estimate the effect of player i 's current report on player j 's flow payoff over the next T periods, where T is large enough that any dependence of j 's flow payoff on i 's report can essentially be

neglected after those T periods. The variable s_i will be defined as the change in the aforementioned effect from its expected value.

The expected payoff of player j over the next T periods depends on j 's previous type θ_j^{-1} , which coincides with the j 's previous report, and i 's current report θ_i . This payoff is evaluated when j 's current type is still unknown. Denote this payoff by

$$B_{k,T}^i(\theta_j^{-1}, \theta_i) = \sum_{t=0}^T (1 - \varepsilon)^t E[u_j^{+t} \mid \theta_i, \theta_j^{-1}],$$

where u_j^{+t} is j 's payoff in t periods after the current one.⁸ We take the value of T to be large enough that the effect of i 's report on j 's payoff after T periods depends only marginally on θ_i .

Claim 1. *For any $\Delta > 0$ and any types θ'_i, θ''_i , and θ_j^{-1} , there exists a number T such that for any $t > T$ we have*

$$|E[u_j^{+t} \mid \theta'_i, \theta_j^{-1}] - E[u_j^{+t} \mid \theta''_i, \theta_j^{-1}]| < \Delta.$$

If players never went on suspension, this claim would follow directly from the convergence of a Markov chain to its limiting distribution η . Since players may go on suspension, with probability depending on their reports, the value of u_j^{+t} may be affected not only by the expectation of the actions that are going to be played but also by the probability of going on suspension. However, the probability of a player's report affecting the possibility of going on suspension in a remote period is of order $O(1/M)$, and is therefore negligible if M is sufficiently large.

We can now define T as the number which satisfies Claim 1 for any Δ less than the difference in the total payoff,

$$[u_1(\theta_1, a(\theta)) + u_2(\theta_2, a(\theta))] - [u_1(\theta_1, a) + u_2(\theta_2, a)],$$

for all profiles $\theta = (\theta_1, \theta_2)$, and all actions $a = a(\theta')$ for $\theta' \neq \theta$. By Assumption I, this difference is positive.

We define $s_k^i(\theta^{-1}, \theta_i)$ as the difference between $B_{k,T}^i(\theta_j^{-1}, \theta_i)$ and its expectation, which is estimated based on i 's report θ_i^{-1} in the previous period:⁹

$$s_k^i(\theta^{-1}, \theta_i) = B_{k,T}^i(\theta_j^{-1}, \theta_i) - E_{\theta_i}[B_{k,T}^i(\theta_j^{-1}, \theta_i) \mid \theta_i^{-1}].$$

We omit T in the notation for s_k^i . The variable p_i is defined as in the i.i.d. case:

$$p_k^i(\theta^{-1}) = \Pr\{s_k^i(\theta^{-1}, \theta_i) > 0\} \cdot E_{\theta_i}[s_k^i(\theta^{-1}, \theta_i) \mid s_k^i(\theta^{-1}, \theta_i) > 0].$$

⁸The formula for $B_{k,T}^i(\theta_j^{-1}, \theta_i)$ includes the possibility that players can be on suspension within the next T periods. Thus, $B_{k,T}^i(\theta_j^{-1}, \theta_i)$ and $s_k^i(\theta^{-1}, \theta_i)$ depend on the number of chips k currently held by player i , in contrast to the case when players had i.i.d. types.

⁹When player $x \in \{i, j\}$ is returning from suspension, we use type θ_x^{-M} that was reported M periods ago, just before the suspension. For the other player $-x$ we use type θ_x^{-1} that was reported in the previous period.

6.2 Efficiency and incentives

In this section, we show that the “Markovian” chip strategies approximate efficient outcomes and are incentive compatible.

Theorem 3. *In any two-player repeated game in which the players’ types are Markov, and the stage game satisfies Assumptions I–III, the efficient payoff can be arbitrarily closely approximated by chip-strategy strict equilibria when the discount factor δ is sufficiently close to 1.*

The transition probabilities between various chip structures in the chip strategies are the same as in the i.i.d. case. Thus, the ergodic distribution over chip structures is also the same as in the i.i.d. case, and the probability of being in an inefficient state, that is, of some player being on suspension, vanishes as $n \rightarrow \infty$.¹⁰ Therefore, players’ payoffs are approximately efficient.

To check that players have incentives to report truthfully, we will show that the value of having one more chip in the Markovian case is very close to the corresponding value in the i.i.d. case when types are distributed according to η . Then we will conclude that the effect of player i ’s report on her payoff is very close to that in the d’Aspremont Gerard-Varet mechanism.

If players report truthfully, the value function of player i , estimated at the beginning of the period, depends on the state, i.e., the number k of chips she holds, and on the previous type profile θ^{-1} . Thus, the value function will be denoted by $V_{k,\theta^{-1}}^i$.¹¹ We first claim that the continuation payoff of player i with k chips depends only in a limited way on the previous type profile θ^{-1} :

Claim 2. *For any k , and two type profiles in the previous period, $\theta^{-1} = \theta$ and $\theta^{-1} = \theta'$, there is a constant $C > 0$, independent of ε, M , and n , such that $|V_{k,\theta}^i - V_{k,\theta'}^i| < C\varepsilon$.*

This claim follows from two facts concerning the prescribed strategies: (a) for any two current type profiles $\theta^{-1} = \theta$ and $\theta^{-1} = \theta'$, the probability that the type profiles t periods from now will coincide tends to 1 at an exponential rate, independent of the discount factor; and (b) for any current number of chips k , the probability of player i having one less (more) chip in the following period (in M periods from now, when $k = 0$ or $2n$) is independent of the previous type profile. More precisely, the probability of having one less (more) chip is $1/3$ ($1/3$) if $1 \leq k \leq 2n - 1$, and is 0 (1) and 1 (0) when $k = 0$ or $2n$, respectively). This probability depends on the current type profile when $1 \leq k \leq 2n - 1$, but in the expectation over all current type profiles, and contingent on any previous type profile, is equal to $1/3$.

Let

$$\bar{V}_k^i = \sum_{\theta^{-1}} \eta(\theta^{-1}) V_{k,\theta^{-1}}^i$$

¹⁰In the Markovian case, the state is characterized not only by the number of chips each player has but also by the type profile θ^{-1} in the previous period, since the transition probabilities depend on this type profile.

¹¹When player $x \in \{i, j\}$ is returning from suspension, the value function depends on θ_x^{-M} (the type profile player x had M periods ago, just before the suspension), and it depends on θ_{-x}^{-1} .

be the expected value of player i with k chips, where the expectations are taken with respect to the limiting distribution η over θ^{-1} . Then the recursive equation for \bar{V}_k^i , for $1 \leq k \leq 2n - 1$ is:

$$\bar{V}_k^i = \varepsilon w^i + (1 - \varepsilon) \frac{1}{3} \bar{V}_{k-1}^i + (1 - \varepsilon) \frac{1}{3} \bar{V}_{k+1}^i + (1 - \varepsilon) \frac{1}{3} \bar{V}_k^i + \varepsilon X_k, \quad (7)$$

where X_k is an expression of order $O(1/M)$. Indeed, if the number of chips held by player i in the next period were independent of the current type profile θ , then (7) would not even have the term εX_k . However, there is a correlation between the number of chips that player i will have in the next period and the current type profile θ . This correlation introduces terms of $V_{k,\theta}^i - V_{k,\theta'}^i$ to expression (7) - compared to the case with i.i.d. types - which are multiplied by some coefficients. By Claim 2, each term $V_{k,\theta}^i - V_{k,\theta'}^i$ is of order $C\varepsilon$, and its multiplying coefficient depends on the probability of players' reports affecting the state, which is of order $O(1/M)$. Hence, term X_k in expression (7) is of order $O(1/M)$.

Therefore, the recursive equation for \bar{V}_k^i , for $1 \leq k \leq 2n - 1$, differs from its i.i.d. analogue, with types distributed according to η , by the term εX_k . Moreover, the expression for \bar{V}_k^i coincides with that in the i.i.d. case for $k = 0, 2n$, since the transitions in the states with $k = 0, 2n$ are independent of types.

We will now estimate player i 's average future benefit from having one more chip:

$$\bar{\Delta}_k^i := \bar{V}_{k+1}^i - \bar{V}_k^i,$$

and show that this term is of order $O(M\varepsilon)$. As in the i.i.d. case, we evaluate the derivative of $\bar{\Delta}_k^i$ with respect to ε by the Implicit Function Theorem, which yields the following system of linear equations:

$$\begin{bmatrix} 1 & 0 & 0 & \cdot & \cdot & \cdot & \cdot \\ -1/3 & 2/3 & -1/3 & \cdot & \cdot & \cdot & \cdot \\ 0 & -1/3 & 2/3 & -1/3 & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & -1/3 & 2/3 & -1/3 \\ \cdot & \cdot & \cdot & \cdot & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \partial \bar{\Delta}_{2n-1}^i / \partial \varepsilon \\ \partial \bar{\Delta}_{2n-2}^i / \partial \varepsilon \\ \cdot \\ \cdot \\ \cdot \\ \partial \bar{\Delta}_0^i / \partial \varepsilon \end{bmatrix} = \begin{bmatrix} Mw_j^i - Mw^i \\ X_{2n-1} - X_{2n-2} \\ X_{2n-2} - X_{2n-3} \\ \cdot \\ \cdot \\ X_2 - X_1 \\ Mw^i - Mv^i \end{bmatrix}.$$

The solution to the system is obtained by Gauss-Jordan elimination method. The difference between this system and that for the i.i.d. case comes from the presence of the expressions X_k , which are of order $O(1/M)$. Thus, the difference between the solutions in the Markov case and those in the i.i.d. case,

$$\partial \Delta_k^i / \partial \varepsilon - \partial \bar{\Delta}_k^i / \partial \varepsilon,$$

is of order $O(n/M)$. Since we assumed that $n \sim M^{3/2}$ (see (6)), the difference $\partial \Delta_k^i / \partial \varepsilon - \partial \bar{\Delta}_k^i / \partial \varepsilon$ is of order $O(M^{1/2})$.

Since terms $V_{k,\theta}^i - V_{k,\theta'}^i$ are of order ε , so are terms $V_{k,\theta}^i - \bar{V}_k^i$, and terms $\Delta_k^i - \bar{\Delta}_k^i$ are of order $\varepsilon O(M^{1/2})$. Thus, by (2) and the fact that terms α_k^i and ϕ_k^i are of order $O(1/M)$, we have that for any k, θ :

$$\alpha_k^i(1 - \varepsilon) [V_{k+1,\theta}^i - V_{k,\theta}^i] = \varepsilon(1 + O(1/M^{1/2})) \text{ and } \phi_k^i(1 - \varepsilon)[V_{k,\theta}^i - V_{k-1,\theta}^i] = \varepsilon(1 + O(1/M^{1/2})).$$

It follows from these formulas that any change $ds_k^i(\theta^{-1}, \theta_i)$ in the value of $s_k^i(\theta^{-1}, \theta_i)$ will affect i 's continuation payoff by $\varepsilon(1 + O(1/M^{1/2}))ds_k^i(\theta^{-1}, \theta_i)$, while in the i.i.d. case any change ds_i in the value of s_i affected i 's continuation payoff by εds_i .

To check players' incentives to report truthfully, we must examine the effect of each player i 's current report $\hat{\theta}'_i$ on her continuation payoff. By referring to the one-stage deviation principle, we can assume that player i will always report truthfully starting from the next period. By inspecting the prescribed strategies, we see that player i 's current report $\hat{\theta}'_i$ affects the following: the current flow payoff, the value of $s_k^i(\theta^{-1}, \hat{\theta}'_i)$ in the current period, and value of $s_{k+1}^i(\hat{\theta}', \theta_i^{+1})$ in the next period, where $\hat{\theta}'$ is the type profile reported in the current period.¹² The former value is affected through the value of $B_{k,T}^i(\theta_j^{-1}, \hat{\theta}'_i)$, and the latter value is affected through the value of $E_{\theta_i^{+1}}[B_{k+1,T}^i(\theta_j, \theta_i^{+1}) \mid \hat{\theta}'_i]$.

In total, the effect of player i 's report $\hat{\theta}'_i$ on her continuation payoff is

$$\begin{aligned} & B_{k,T}^i(\theta_j^{-1}, \hat{\theta}'_i)(\varepsilon + O(\varepsilon/M^{1/2})) - (1 - \varepsilon)E_{\theta_j}E_{\theta_i^{+1}}[B_{k+1,T}^i(\theta_j, \theta_i^{+1}) \mid \hat{\theta}'_i, \theta_j^{-1}](\varepsilon + O(\varepsilon/M^{1/2})) \\ & = \varepsilon[B_{k,T}^i(\theta_j^{-1}, \hat{\theta}'_i) - (1 - \varepsilon)E_{\theta_j}E_{\theta_i^{+1}}[B_{k+1,T}^i(\theta_j, \theta_i^{+1}) \mid \hat{\theta}'_i, \theta_j^{-1}]] + O(\varepsilon/M^{1/2}). \end{aligned}$$

The term $O(\varepsilon/M^{1/2})$ does not affect incentives, and hence can be omitted. Recalling the definition of $B_{k,T}^i(\theta_j^{-1}, \hat{\theta}'_i)$, we obtain

$$\begin{aligned} & B_{k,T}^i(\theta_j^{-1}, \hat{\theta}'_i) - (1 - \varepsilon)E_{\theta_j}E_{\theta_i^{+1}}[B_{k+1,T}^i(\theta_j, \theta_i^{+1}) \mid \hat{\theta}'_i, \theta_j^{-1}] \\ & = \sum_{t=0}^T (1 - \varepsilon)^t E[u_j^{+t} \mid \hat{\theta}'_i, \theta_j^{-1}] - (1 - \varepsilon) \sum_{t=0}^T (1 - \varepsilon)^t E[u_j^{+t+1} \mid \hat{\theta}'_i, \theta_j^{-1}] \\ & = E[u_j \mid \hat{\theta}'_i, \theta_j^{-1}] - (1 - \varepsilon)^{T+1} E[u_j^{+T+1} \mid \hat{\theta}'_i, \theta_j^{-1}]. \end{aligned} \tag{8}$$

The first term in (8) is equal to

$$E_{\theta_j}[u_j(\theta_j, a(\hat{\theta}'_i, \theta_j)) \mid \hat{\theta}'_i, \theta_j^{-1}],$$

and together with the effect of player i 's report on her current flow payoff, it gives her incentives to maximize the total payoff; while the second term in (8) depends on $\hat{\theta}'_i$ up to some value less than Δ , and therefore is inessential to player i 's incentives.

¹²In the next period, the number of chips held by player i may differ from k . We denote that number by k^{+1} .

7 Efficient chip strategies for games with any number of players

We will modify the chip strategies constructed previously for two-player games to accommodate games with any number of players. Under some assumptions on the stage game, the modified chip strategies will also constitute an (approximately) efficient equilibrium. The chip strategies in games with more than two players prescribe several stages. In each stage, there is a non-empty subset $R \subset I$ of active players, and the total payoff of set R is maximized. Over time, some players may become active and other players may have to leave set R , by going on suspension. As a result, the stage of the game may change. At the beginning, all the players are active.

In this section, we will describe chip strategies only in the stage in which all players are active, and postpone describing strategies in other stages (in which some players are on suspension) to Section 7.1. We modify the strategies from Section 4 as follows:

Before the game begins, every player is given n chips. Then one player is selected at random, and taken away a chip (which is shared equally by the remaining $I - 1$ players). In each period, players exchange chips, and given that all players are still active (i.e., no player goes on suspension), exactly one player will have fewer than n chips. All her missing chips are always shared equally among the other (active) players. If it happens at the end of some period, that each player has exactly n chips, then a player is selected randomly from the set of all players in an equal-chance lottery, and one chip is taken away from the selected player.

We denote the state in which player i has $0 < k < n$ chips by Ω_k^i . We modify (C1)–(C3) in the definition of the chip strategies from Section 4 to (K1)–(K3) as follows:

(K1) In state Ω_k^i , player i obtains a chip from the opponents (a fraction of $1/(I - 1)$ from each opponent) with probability $\alpha_k^i s_i \mathbf{1}_{\{s_i > 0\}}$, and player i gives the opponents a chip (a fraction of $1/(I - 1)$ to each opponent) with probability $-\phi_k^i s_i \mathbf{1}_{\{s_i \leq 0\}}$, where α_k^i and ϕ_k^i are defined as in the two-player case;

(K2) In state Ω_k^i , player i obtains a chip from the opponents with probability $-\sum_{j \neq i} \phi_{i,k}^j s_j \mathbf{1}_{\{s_j \leq 0\}}$, and player i gives the opponents a chip with probability $\sum_{j \neq i} \alpha_{i,k}^j s_j \mathbf{1}_{\{s_j > 0\}}$, where $\alpha_{i,k}^j$ and $\phi_{i,k}^j$ will be defined shortly;

(K3) In state Ω_k^i , player i gives the opponents a chip with probability

$$\frac{1}{3} - \phi_k^i p_i - \sum_{j \neq i} \alpha_{i,k}^j p_j,$$

independently of the messages sent in the current period; and player i obtains from the opponents a chip with probability

$$\frac{1}{3} - \alpha_k^i p_i - \sum_{j \neq i} \phi_{i,k}^j p_j,$$

independent of the messages sent in the current period.

Part (D) is as in Section 4, and part (E) will be discussed in the next section.

Intuitively, a player with fewer than n chips is provided incentives to maximize the total payoff by a higher chance of obtaining a chip versus being taken a chip away, as in the two-player case. In turn, players with more than n chips are provided incentives (player by player) by a higher chance that the player with fewer than n chips will be taken away one more chip (versus the chance that she will obtain a chip back).

The coefficients α_k^i and ϕ_k^i for $0 < k < n$, and the coefficients $\alpha_{i,k}^j$ and $\phi_{i,k}^j$ for $j \neq i$ and $0 < k < n$, are defined in such a way that equation (2) and the following equations are satisfied:

$$\alpha_{i,k}^j(1 - \varepsilon)[W_{i,k}^j - W_{i,k+1}^j] = \varepsilon \text{ and } \phi_{i,k}^j(1 - \varepsilon) [W_{i,k-1}^j - W_{i,k}^j] = \varepsilon,$$

where $W_{i,k}^j$ stands for the continuation payoff of player j in state Ω_k^i . These two last equations guarantee that in state Ω_k^i player j maximizes the total (across all players) stage-game payoff. As in Section 4, the V 's and the W 's are determined by recursive expressions: equations (3)–(5) for the V 's, the corresponding equations for the W 's, and the equation

$$W_{j,n}^i = V_n^i = \frac{1}{I} \left[V_{n-1}^i + \sum_{j \neq i} W_{j,n-1}^i \right].$$

This last equation must hold because our strategies prescribe that when each player has exactly n chips, a player is selected randomly from the set of all players in an equal-chance lottery, and a chip is taken away from the selected player.

These chip strategies are approximately efficient when n is sufficiently large, but the incentive constraints are satisfied only under the following additional assumption:

Assumption IV: For any $i, j \in \{1, \dots, I\}$ and any R such that $i, j \in R$,

$$\frac{1}{|R|} \left(\sum_{R \ni l \neq i} w_{R-\{l\}}^i + v_{R-\{i\}}^i \right) < w_{R-\{j\}}^i.$$

The assumption says that any player i prefers the suspension of any player j other than herself to the suspension of a randomly chosen player (including herself). Assumption IV is satisfied in many applications, including all anonymous stage games. However, Assumption IV is strong. In particular, it implies that the payoff of any player i when some player $l \neq i$ is on suspension cannot be much larger than player i 's payoff when another player $j \neq i, l$ is on suspension. The assumption seems necessary for the modified construction. Intuitively, if player j is currently holding fewer than n chips, then player i must be willing to sacrifice some flow payoff to increase the chance that player j is taken away one more chip, as opposed to the chance that player j will obtain one chip back. The former event will increase the chance that player j will go on suspension, and the latter event will increase the chance that some random player (including player i) will go on suspension. This is exactly what is guaranteed by Assumption IV.

Theorem 4. *If the players' types are i.i.d. (or more generally Markov) and the stage game satisfies Assumptions I–IV, then the efficient payoff can be approximated by chip-strategy strict equilibria when the discount factor δ approaches 1.*

Of course, we still need to specify the strategies at histories in which some players are on suspension. Such histories, except those in which only one player is active, are specific for the games with more than two players. We discuss “the play on suspension” in the following section. Once the strategies at histories in which some players are on suspension are specified, it will become clear that the proof of Theorem 4 closely mimics the proofs of Theorems 2 and 3. So, we will omit the details of this proof.

Remark 1. Of course, there are many other ways of defining transitions between various states. They need to have only the properties indicated in boldface in Section 5. Perhaps some alternative transitions would allow to prove Theorem 4 under a different set of assumptions, or even under a somewhat relaxed version of Assumption IV. We have not explored all of the alternative ways. We conjecture, however, that it would be necessary to assume a bound on the payoffs of a player i when the opponent j is suspended, by a function of player i 's payoff when the opponent l is suspended.

7.1 Play on suspension

To complete the analysis, we need to specify the play at stages when some players are on suspension. First, notice that our analysis of the two-player case is valid if we assume that players go on suspension not for a deterministic number of periods, M , but for a random number of periods. That is, in every period, a player on suspension is allowed to come back and become active with probability μ , such that the expected length of suspension is

$$M = \sum_{t=1}^{\infty} t\mu(1-\mu)^{t-1} = \frac{1}{\mu}. \quad (9)$$

We define the repeated-game strategies as follows:

For each number $J < I$ of active players, denote the values $n[J]$ and $M[J]$ as the number of chips per player and the expected length of suspension at any stage with J active players. If all the players are active, we have $n[I] = n$ and $M[I] = M$.

Recall that Ω_k^i denotes the state in which all players are active and i is the player with $k < n$ chips. Once player i_1 goes on suspension, the play moves to state $\Omega_{n[I-1]-1}^{i_1, i}$ in which player i , who is randomly chosen from the players other than i_1 , gives away one chip. In states $\Omega_k^{i_1, i}$ with $0 < k < n[I-1]$, players maximize the total payoff of the players other than i_1 , and the transitions between states are defined as in the repeated game in which there are $I-1$ players and they are all active. In any state $\Omega_k^{i_1, i}$, player i_1 can return from suspension, in which case the play moves to state $\Omega_1^{i_1}$, that is, one chip is returned to player i_1 (so she holds 1 chip). What happened in states $\Omega_k^{i_1, i}$ becomes irrelevant: If player i_1 goes on suspension again, then the play moves to state $\Omega_{n[I-1]-1}^{i_1, i}$ for the player i who is randomly selected from the players other than i_1 .

If player i_1 is on suspension and the game reaches the stage in which another player i_2 goes on suspension (for an expected length of $M[I-1]$ periods), a player i is randomly selected from the players other than i_1 and i_2 , and the play moves to state $\Omega_{n[I-2]-1}^{i_1, i_2, i}$. In state $\Omega_{n[I-2]-1}^{i_1, i_2, i}$, the players maximize the total payoff of

the active players, that is, all players but i_1 and i_2 , and the transitions between states are defined as in the repeated game in which there are $I - 2$ players and they are all active. In each period, either of the players i_1, i_2 may return from suspension. We impose the following condition: if a player returns from suspension, all the players who went on suspension after her, return to the game as well. That is, if player i_1 returns, then player i_2 returns to the game as well, and the play moves to state $\Omega_1^{i_1}$. One chip is returned to player i_1 , and what happened after the period in which i_1 went on suspension becomes irrelevant. If player i_2 returns from suspension before player i_1 , then the play moves to state $\Omega_1^{i_1, i_2}$. That is, one chip is returned to player i_2 , and what happened after the time that player i_2 went on suspension becomes irrelevant.

More generally, for any sequence of players i_1, i_2, \dots, i_l on suspension, the total payoff of the active players is maximized in states $\Omega_k^{i_1, i_2, \dots, i_l, i}$ with $0 < k < n[I - l]$, and the transitions between states are defined as in the repeated game in which there are $I - l$ players and they are all active.¹³ If an active player i_{l+1} goes on suspension, the play moves to state $\Omega_{n[I-l]-1}^{i_1, i_2, \dots, i_l, i_{l+1}, i}$ in which a randomly selected player $i \neq i_1, i_2, \dots, i_l, i_{l+1}$ begins with $n[I - l - 1] - 1$ chips. If player $i_m, m = 1, \dots, l$, returns from suspension, then the players i_{m+1}, \dots, i_l return as well, and the play moves to state $\Omega_1^{i_1, i_2, \dots, i_m}$.

We will now present a concrete example of chip strategies in a game with three players. The outcome implemented by these strategies will be inefficient. The purpose of this example is rather to exhibit the simplest possible strategies representing our construction. In particular, the length of suspension periods will be deterministic, and the rules governing the chip transition will also be deterministic, contingent on the choice of player whose report determines the chip transition.

Example. Consider the following model of favor exchange. In each period, each player can provide a favor to the two other players, whose benefit from receiving the favor is 1 per player. The cost to the provider is i.i.d., taking values $2 + \vartheta$ or ϑ with a fifty-fifty chance. If $\vartheta < 2$, the efficient action of each player is to provide the favor only if the cost is equal to ϑ . And if $\vartheta > 0$, no player has ever an incentive to provide a favor in the stage game.

Consider now the following chip strategies. Each player begins with $n = n[3] = 2$ chips, and each player is initially active. Then, one player is selected at random, and taken away a chip. Next, players truthfully announce their costs, and play the efficient actions. After that, one of them is selected randomly. If this is the player with one chip, and she reported the cost of ϑ , she obtains a second chip; if she reported the cost of $2 + \vartheta$, she is taken away the only her chip, and goes on suspension. If this is the player with more than one chip, and she reported the cost of $2 + \vartheta$, the player with one chip obtains a second chip; if she reported the cost of ϑ , the player with one chip is taken away her only chip, and goes on suspension.

The suspension of any player lasts $M[3] = 1$ period, in which each active player has $n[2] = 1$ chip. In this state, one of the two active players is selected at random, and taken away a chip. This means that the

¹³In particular, if player i is the only active player, then players take the actions that maximize i 's payoff, given i 's reported type.

selected player goes on immediate suspension. The suspension lasts $M[2] = 1$ period. In other words, when one of the three players goes on suspension because of losing both chips, another player (selected randomly) also goes on suspension as well. And after one period of suspension both players return, and the game is continued with the player who went on suspension first having one chip. Players on suspension provide favors, independently of their costs, while the only active player provides no favors, also independently of the cost being equal to $2 + \vartheta$ or ϑ .

In this game, $w_R^i = 1 - 0.5\vartheta$ when $|R| = 2$ and $i \in R$, and $v_R^i = -\vartheta$ when $|R| = 2$ and $i \notin R$. That is, the expected flow payoff of a player who is not on suspension is $1 - 0.5\vartheta$, and it is only $-\vartheta$ for the player on suspension. By symmetry of payoffs and strategies, $W_{i,k}^j$ - the continuation payoff of player $j \neq i$ when player i has $k = 0, 1, 2$ chips - is the same for all i and j , and will be denoted by W_k ; similarly, V_k^i - the continuation payoff of player i when player i has $k = 0, 1, 2$ chips - is the same for all i , and will be denoted by V_k . Both W_0 and V_0 refer to the states such that $|R| = 2$. Notice that $W_2 = V_2$, and will be denoted by U .

Then,

$$\begin{aligned} W_0 &= (1 - \delta) \cdot (1 - 0.5\vartheta) + \delta W_1 \\ W_1 &= (1 - \delta) \cdot (1 - 0.5\vartheta) + \frac{1}{2}\delta U + \frac{1}{2}\delta W_0 \\ U &= \frac{1}{3}V_1 + \frac{2}{3}W_1 \\ V_1 &= (1 - \delta) \cdot (1 - 0.5\vartheta) + \frac{1}{2}\delta U + \frac{1}{2}\delta V_0 \\ V_0 &= (1 - \delta) \cdot (-\vartheta) + \delta V_1 \end{aligned}$$

These equations determine W_k 's and V_k 's.

To show that the prescribed strategies constitute an equilibrium, we must check that players have incentives to report their types truthfully when one of the players has one chip. This means that

$$(1 - \delta)\vartheta \leq \frac{1}{3}\delta(U - V_0) \leq (1 - \delta)(2 + \vartheta)$$

for the player who has 1 chip, and that

$$(1 - \delta)\vartheta \leq \frac{1}{3}\delta(W_0 - U) \leq (1 - \delta)(2 + \vartheta)$$

for the opponents of that player. These inequalities are satisfied, for example, for $\delta = 0.9$ and $\vartheta = 0.1$.

It seems useful to explain some relation between our equilibria and the equilibria from some closely related papers. Our construction is in a way similar to the constructions used in Fudenberg et al. (1994) (henceforth FLM), and Hörner et al. (2015) (henceforth HTV). In those papers, players always maximize a certain weighted average of their payoffs. On the equilibrium paths of our equilibria, players maximize the total payoff of a subset of players. So, they maximize only specific weighted averages. In addition, the weighted averages of players' payoffs which are maximized in FLM and HTV keep actively changing over

time, while chip strategies use only some fixed set of weighted averages. This feature allows for providing a more detailed description of players' behavior compared to these two other papers, since one only needs to know the set of active players to prescribe actions (reports) for the current period.

The changes in weights are used in FLM and HTV to provide players incentives to play the prescribed actions by using player-specific punishments. This requires assuming that the set of stage-game payoff vectors is full dimensional. The corresponding role is played in our setting by Assumption IV. Actually, Assumption IV, together with the existence of a bad equilibrium imply that the set of stage-game payoff vectors is full dimensional. Similarly, Assumption IV is related to the nonempty interior assumption in Escobar and Toikka (2011), although it is used in a slightly different way. Escobar and Toikka use nonempty interior to impose player-specific punishments only off equilibrium path. In chip strategies, the inefficient stages are used to deter players from misreporting their types and occur on the equilibrium path.

It should be almost clear that we can prove Theorem 4 by analogous arguments to those used to prove Theorems 2 and 3. Only some additional comments seem important and necessary:

1) From the perspective of a state with $I - l$ active players, the probability of returning to a state with a higher number of active players, that is, when one of the players i_1, i_2, \dots, i_l comes back from suspension, can be thought of as additional discounting. This probability, so the additional discounting, vanish when the discounting in the repeated game vanishes, and when all expected lengths of suspension $M, M[I-1], \dots, M[I-l+1]$ tend to infinity at an appropriate rate. More precisely, we require that at any stage with x active players, the ratios $\frac{n[x]}{M[x]}$, $\frac{M^2[x]}{n[x]}$, and $\frac{M[x]}{n[x-1]}$ tend to infinity. Under these conditions, we can use an inductive argument, backward with respect to the number of players on suspension, and show that the expected payoff of every player i in every state in which all the players are active converges to the efficient payoff w^i .

2) Among any set of active $I - l$ players, there is exactly one player i who holds fewer chips than anyone else, and is thus more likely to go on suspension. Our assumptions guarantee that each of active $I - l$ players assigns a positive value to having more chips. More specifically, Assumption II guarantees that player i does not want to go on suspension, and Assumption IV guarantees that any other active player $j \neq i$ prefers player i to go on suspension to some random player (including player j) going on suspension. These two assumptions enable us to find an appropriate probability in which a player's report affects the state such that the player has incentives to report truthfully. In addition, players have incentives for not deviating from the prescribed actions, because the continuation payoff of each player, whether on suspension or not, is arbitrarily close to an efficient payoff. And Assumption III enables us to discipline players by the perspective of playing the bad equilibrium.

3) Finally, we need to justify the recursive formulas for V_n^i and $W_{j,n}^i$ for the game in which all the players are active. For the two-player game, those formulas were obtained under the assumption that when a player is on suspension, player i obtains the stage-game payoff of v^i or w_j^i (depending on whether i or $j \neq i$ is on suspension) for M consecutive periods. Note that this is true in approximation for any number of players,

because the probability of any other player going on suspension when one player is already on suspension tends to 0 when $n \rightarrow \infty$.

8 Advantages of chip strategies

One advantage of chip strategies over other strategies used for similar purposes in the existing literature (such as review strategies) is that players condition their actions on only a simple statistics of the past play, especially when the number of players is small. In the case of two players, for example, if the play is in the cooperation phase (i.e., no player is on suspension), they condition only on the number of chips each player has, and—when types are Markov—the type profile reported in the previous period. If a player is on suspension, they condition on the number of periods in which that player has already been suspended.

In particular, chip strategies depend only minimally (and indirectly) on the space of types and actions. Therefore, they seem particularly attractive in games with large numbers of types and actions. Indeed, repeated-game strategies typically prescribe different actions for different types, and must give players incentives to play such actions. For example, review strategies do so by dividing the time horizon into long blocks, and performing a test at the end of each block to check whether the frequency of actions taken within that block is close to the frequency that would be expected if the players had taken the prescribed actions. If the numbers of both types and actions are large, then review strategies require performing a large number of frequency tests, making chip strategies to be a relatively simple and attractive alternative.

9 References

Abdulkadiroğlu, Atila and Kyle Bagwell (2012). “The Optimal Chips Mechanism in a Model of Favors,” mimeo.

Abdulkadiroğlu, Atila and Kyle Bagwell (2013). “Trust, Reciprocity and Favors in Cooperative Relationships,” *American Economic Journal: Microeconomics*, 5(2), 213 – 259.

Aoyagi, Masaki (2003). “Bid Rotation and Collusion in Repeated Auctions,” *Journal of Economic Theory*, 112, 79–106.

Aoyagi, Masaki (2007), “Efficient Collusion in Repeated Auctions with Communication,” *Journal of Economic Theory*, 134, 61–92.

Arrow, Kenneth J. (1979). “The Property Rights Doctrine and Demand Revelation Under Incomplete Information,” in *Economics and Human Welfare*, ed. by M. Boskin, New York, Academic Press.

Athey, Susan and Kyle Bagwell (2008). “Collusion with Persistent Cost Shocks,” *Econometrica*, 76(3), 493–540.

Athey, Susan, Kyle Bagwell, and Chris W. Sanchirico (2004). “Collusion and Price Rigidity,” *Review of Economic Studies*, 71(2), 317–349.

- Athey, Susan and Ilya Segal (2013). “An Efficient Dynamic Mechanism,” *Econometrica* 81(6), 2463–2485.
- Blume, Andreas and Paul Heidhues (2002). “Modelling Tacit Collusion in Repeated Auctions,” *Journal of Institutional and Theoretical Economics*, 164(1), 163–184.
- d’Aspremont, Claude and Louis-André Gérard-Varet (1979). “Incentives and Incomplete information,” *Journal of Public Economics* 11, 25–45.
- Escobar, Juan F. and Juuso Toikka (2013). “Efficiency in Games with Markovian Private Information,” *Econometrica* 81, 1887–1934.
- Fudenberg, Drew, David Levine, and Eric Maskin (1994). “The Folk Theorem with Imperfect Public Information,” *Econometrica*, 62(5), 997–1039.
- Hauser, Christine and Hugo Hopenhayn (2008). “Trading Favors: Optimal Exchange and Forgiveness,” mimeo.
- Hörner, Johannes, Satoru Takahashi, and Nicolas Vieille (2015). “Truthful Equilibria in Dynamic Bayesian Games,” *Econometrica*, 83, 1795–1848.
- Jackson, Matthew O. and Hugo F. Sonnenschein (2007). “Overcoming Incentive Constraints by Linking Decisions,” *Econometrica*, 75(1), 241–258.
- Mailath, George J. and Larry Samuelson (2006). *Repeated Games and Reputations*, Oxford University Press, Oxford.
- Möbius, Markus M. (2001). “Trading Favors,” mimeo.
- Olszewski, Wojciech and Mikhail Safonov (2017). “Efficient Cooperation by Exchanging Favors,” mimeo.
- Rachmilevitch, Shiran (2013). “Endogenous Bid Rotation in Repeated Auctions,” *Journal of Economic Theory*, 148, 1714 – 1725.
- Radner, Roy (1985). “Repeated Principal-Agent Games with Discounting,” *Econometrica*, 53(5), 1173–1198.
- Shiryaev, Albert N. (1996) *Probability*, 2nd edition, Springer-Verlag, New York, Inc. 1996.
- Skrzypacz, Andrzej and Hugo Hopenhayn (2004). “Tacit Collusion in Repeated Auctions,” *Journal of Economic Theory*, 114(1), 153–69.
- Spulber, Daniel F. (1995). “Bertrand Competition When Rivals’ Costs Are Unknown,” *Journal of Industrial Economics*, 43(1), 1–11.
- Wolitzky, Alexander (2015). “Communication with Tokens in Repeated Games on Networks,” *Theoretical Economics*, 10, 67–101.