

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Optimal stealing time

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1571711> since 2018-11-06T14:58:59Z

Published version:

DOI:10.1007/s11238-015-9507-y

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Optimal Stealing Time

Andrea Gallice*

University of Torino and Collegio Carlo Alberto

The published version of this paper can be downloaded at the following link:

<https://link.springer.com/article/10.1007/s11238-015-9507-y>

Abstract

We study a dynamic game in which players can steal parts of a homogeneous and perfectly divisible pie from each other. The effectiveness of a player's theft is a random function which is stochastically increasing in the share of the pie the agent currently owns. We show how the incentives to preempt or to follow the rivals change with the number of players involved in the game and investigate the conditions that lead to the occurrence of symmetric or asymmetric equilibria.

Keywords: Stealing; Stochastic games; Optimal timing; Pie allocation

1 Introduction

In this paper, we introduce and study what we call the stealing game. A stealing game is a dynamic game in which a number of agents steal portions of a

*Address: ESOMAS Department, University of Torino, Corso Unione Sovietica 218bis, 10134, Torino, Italy. Email: andrea.gallice@unito.it.

homogeneous and perfectly divisible pie from each other. The portion of the pie that a player can steal is stochastic. However, the expected value of this random variable is increasing in the agent’s current holdings such that larger players are able on average to steal larger portions. Within such a framework, agents must decide when and whom to rob, with the goal of finishing the game as the leader, i.e., the player who holds the largest share of the pie.¹

Our primary goal is to solve for the optimal timing strategies of the agents. We want to find the best moment for a player to behave aggressively and steal part of the pie owned by his rivals. Such a decision is affected by an intuitive trade-off between preempting or postponing one’s move. A player who moves as soon as possible eliminates the possibility of being preempted, but he is then forced to passively suffer the potential retaliation of those who waited. On the other hand, a player who postpones his move can observe the new state of the world and react optimally. However, the agent faces the risk of being preempted and robbed by a rival, in which case his market share goes down as does the expected effectiveness of his stealing attempt.

We characterize the pure strategy equilibria of the stealing game under different specifications for the number of players, the duration of the game, and the number of stealing possibilities players are endowed with. We start by explicitly solving a two-period stealing game in which players have a single stealing opportunity. Despite its simplicity, this setting highlights the strategic peculiarities of the game and shows how the above-mentioned trade-off has different solutions depending on the number of participants. No player postpones his move in a

¹As an example of a situation that matches some of the key features of the game, consider the case of electoral competition among political candidates. By campaigning on specific topics, a candidate may target a particular opponent and thus “steal” a portion of his voters. Moreover, larger players (i.e., candidates with many supporters) are usually able to raise more funds, so they can afford more expensive campaigns, which are in turn expected to be more effective.

two-player game. A three-player game displays multiple equilibria and, in some of them, all agents postpone their moves. Finally, when the number of players is larger than three, we show that the number of preempting equilibria is strictly larger than the number of postponing equilibria and that asymmetric equilibria may also occur. We then generalize some of these results to a setting in which n players have K stealing opportunities in a stealing game that lasts for $T > K$ periods.

The paper is organized as follows: Section 2 reviews the relevant literature. Section 3 formally introduces the stealing game. Section 4 defines the equilibria of the game when players have a single stealing opportunity and there are only two periods. Section 5 generalizes the results, and Section 6 concludes.

2 Literature review

In terms of approach, modelling strategy, and topic of investigation, the stealing game has ties with various strands of the literature. The game is a timing game, i.e., a game in which agents must decide when to move (in our specific case, when to use their stealing attempts). The stealing game actually shares something in common with different archetypes of timing games. The two-player game belongs, in fact, to the class of preemption games. These are games in which it is better to act before one's rivals; famous examples are the Stackelberg quantity game (Von Stackelberg 1934) and the centipede game (Rosenthal 1981). On the other hand, the game with more than two players displays some features that are typical of a war of attrition (Maynard Smith 1974), a strategic situation in which preempting the others is disadvantageous.²

²More recent literature on timing games has focused on generalizing former results (Bulow and Klemperer 1999), in providing a unified framework to study preemption games and wars of attrition (Park and Smith 2008), or in experimentally testing some of the theoretical results

The stealing game can also be seen as a dynamic contest where agents compete over different periods with the goal of winning a final prize. Within the rich literature on dynamic contests (see Konrad 2009 for a review), the framework recently discussed in a paper by Sela and Erez (2013) shares some similarities with our stealing game. The authors study a specific form of repeated competition where, in any given period, a firm wins a contest against its rival with a probability that is positively affected by the firm's relative allocation of a finite resource. Firms are budget-constrained: they start with a given budget, and this budget is progressively eroded by the allocation they implement in each period. The game thus differ in a number of dimensions with respect to a stealing game such as the payoff structure (a prize in each period versus a unique final prize in our game), the number of players involved ($n = 2$ versus $n \geq 2$ in our game), and the way the agents compete (by allocating resources versus by stealing resources in our game). Nevertheless, an important feature that characterizes both games is the idea that players are budget-constrained (in our case, with respect to the number of stealing possibilities) and thus must choose the timing profile of their actions as well as the fact that, in any given period, the outcome of the interactions among agents is stochastic.

The idea that players compete by allocating finite resources across periods is also reminiscent of the Colonel Blotto game introduced by Borel (1921). In this game, two contestants must simultaneously deploy their armies over various battlefields, and in every battlefield, victory goes to the agent who positions the greater force. The winner of the game is the agent who wins in the majority of battlefields. Indeed, the basic structure of the stealing game is a specific version of a Blotto game where players can deploy at most one unit (i.e., one stealing possibility) over a subset of battlefields/periods. Our game is then enriched by

(Brunnermeier and Morgan 2010).

other elements, such as the presence of more than two players and the positive relationship between a player’s strength and the expected effectiveness of his move.

Concerning the last point, Rinott, Scarsini, and Yu (2012) introduce and study a gladiator game. The game is a stochastic version of the Blotto game where the coaches of two teams of gladiators must decide how to allocate a finite amount of “total strength” within their teams. Gladiators are then involved in a sequence of one-to-one fights. In any given fight, the probability of winning is a probabilistic function of the fighters’ strength. At the end of each fight, the winner recovers his initial strength and remains to fight a new challenger. Thus, some stochastic elements in the determination of the winner as well as a positive relationship between a player’s current strength and his probability of winning are characteristics similar to our stealing game.

Finally, and partly moving to a different strand of the literature, Dubovik and Parakhonyak (2014) study a dynamic model of targeted competition (i.e., a model in which a player can compete/fight against a specific chosen rival). More precisely, three drug cartels compete over three markets, where each market is served by a different couple of cartels. Each cartel can allocate resources to fight the rival in any of the markets where he operates, and the amount of damage that a cartel can inflict on a rival is positively related to the cartel’s local strength, as measured by its manpower. There are thus some important similarities to the approach that we adopt in modeling the stealing game. In fact, the stealing game also provides a model of targeted competition (whenever $n \geq 3$, each player must choose not only when to steal but also from whom). Moreover, our model also features a positive relationship between the current strength of a player and his expected ability to damage a rival. On the other hand, these two models differ in a number of ways. For instance, in our game, the aforementioned relationship is

stochastic rather than deterministic, all players are active in a common market, agents do not accrue payoffs over time, and the analysis is not restricted to a situation with three players.

3 The stealing game

The stealing game is a discrete-time stochastic dynamic game in which $n \geq 2$ risk-neutral players compete for the possession of a perfectly divisible resource the size of which is constant and normalized to 1. Let $\pi_i^t \in [0, 1]$ be the share of the resource that agent $i \in \{1, \dots, n\}$ holds at time $t \in \{1, \dots, T + 1\}$ where $T \geq 2$ is finite and common knowledge. The vector $\pi^t = (\pi_1^t, \dots, \pi_n^t)$ such that $\sum_i \pi_i^t = 1$ thus defines the allocation at time t with $\pi^1 = (\frac{1}{n}, \dots, \frac{1}{n})$.

The goal of the players is to be the largest shareholder at the moment the game is over (i.e., at $t = T + 1$). Throughout the game, the only way in which an agent can increase his holdings is to steal part of the resource from someone else. Each agent is endowed with $K < T$ stealing opportunities. A player's problem consists of deciding when to use these opportunities (agents can use at most one stealing opportunity per period) and which opponent to target (agents can steal from a single rival).³ The vector $k^t = (k_1^t, \dots, k_n^t)$ with $k_i^t \in \{0, \dots, K\}$ describes players' remaining stealing opportunities at the beginning of period t .

The state of the game at time t is thus defined by $\theta^t = (\pi^t, k^t)$. In any period $t \in \{1, \dots, T\}$, agents first observe θ^t and then simultaneously choose whether to

³There are a number of things to notice here. First, we set $K < T$ because we are interested in studying a situation in which stealing opportunities are a scarce resource, and players must decide when to use them. Second, an agent can freely change the rival he targets across periods: agent a may steal from b in a certain period and then from c in a subsequent period. Finally, we assume for simplicity that there are no explicit monetary costs associated with the act of stealing. Such an assumption implies little loss of generality since all the results would remain valid as long as stealing costs do not exceed a certain threshold.

remain inactive (action $a_i^t = \emptyset$) or steal from a specific opponent j (action $a_i^t = j$). Obviously, an agent who runs out of stealing opportunities must necessarily remain inactive. We indicate with A_i^t a player's action space and with $a^t = (a_1^t, \dots, a_n^t)$ an action profile.

Whenever an agent plays $a_i^t = j$, the *maximal* amount $s_i^t \in [0, 1]$ that agent i can steal from j is determined by the realization of the random variable S_i^t . Let \bar{s}_i^t denote the expected value of S_i^t . The following assumption states that on average larger players are better thieves, i.e., players whose stealing attempts are expected to be more effective.

Assumption 1 $\bar{s}_i^t = f(\pi_i^t)$ with $f(0) = 0$ and $f'(\cdot) > 0$.

Clearly, there may be cases of “excess demand”, i.e., situations in which one or more players simultaneously steal from agent j but j 's holdings are not enough to satisfy aggregate demand. More formally, $\sum_{l: a_l^t = j} s_l^t > \pi_j^t$. Whenever such an event occurs, we assume that thieves obtain a share that is proportional to the strength of their stealing attempts. We can thus define the *actual* amount $y_i^t \leq s_i^t$ that agent i manages to steal from j as follows

$$y_i^t = \min \left\{ s_i^t, \frac{s_i^t}{\sum_{l: a_l^t = j} s_l^t} \pi_j^t \right\} \quad (1)$$

Agents' payoffs are determined by the final allocation of the resource. More precisely, the player who, in period $T + 1$, holds the largest share gets a prize of size 1. The others get zero. If there is more than one market leader, the prize is equally shared among the winners.

We follow the standard practice in the stochastic games literature (see for instance Maskin and Tirole 2001) and focus on Markov strategies. Let $h_t = (a^1, \dots, a^{t-1})$ be the history of the game at the beginning of period t , i.e., the sequence of actions chosen up to period $t - 1$. A Markov strategy depends only

on the current state θ^t and not on the entire history of play h_t . Formally, it is given by $\sigma_i = (\sigma_i^1, \dots, \sigma_i^T)$ where σ_i^t maps the current state into actions, i.e., $\sigma_i^t : \Theta \rightarrow A_i^t$. Notice that we only consider pure strategies as every σ_i^t selects a specific action in A_i^t and does not involve any randomization. We indicate with $\sigma = (\sigma_1, \dots, \sigma_n)$ a profile of pure Markov strategies.

We use as a solution concept the notion of Markov perfect equilibrium (MPE). A MPE is a subgame perfect equilibrium in which all players use Markov strategies (see again Maskin and Tirole 2001). Let $\bar{u}_i(\sigma)$ indicate the agent's expected payoff given the strategy profile σ . A profile $\hat{\sigma} = (\hat{\sigma}_1, \dots, \hat{\sigma}_n)$ is a pure strategy MPE if for any t , h_t , and i

$$\bar{u}_i(\hat{\sigma}_i, \hat{\sigma}_{-i} \mid h_t) \geq \bar{u}_i(\sigma_i, \hat{\sigma}_{-i} \mid h_t) \text{ for all } \sigma_i. \quad (2)$$

3.1 Some preliminary results

The following lemma, whose proof is trivial and is therefore omitted, reduces the set of strategies that can be part of a MPE. It states that in any equilibrium, all agents use all their stealing opportunities.

Lemma 1 *Let the strategy profile $\hat{\sigma} = (\hat{\sigma}_1, \dots, \hat{\sigma}_n)$ be a Markov perfect equilibrium. Then, $\hat{\sigma}_i$ is such that $k_i^{T+1} = 0$ for any $i \in N$.*

Lemma 2 defines instead the relationship between an agent's current holdings (π_i^t) and the expected value of his loot (\bar{y}_i^t). It states that \bar{y}_i^t is strictly increasing in π_i^t . The result immediately follows from Assumption 1 and the definition of y_i^t (see expression 1).

Lemma 2 $\bar{y}_i^t = g(\pi_i^t)$ with $g(0) = 0$ and $g'(\cdot) > 0$.

In the remaining of the paper, we will also extensively use the notion of a “circle” of players.

Definition 3 *Given any set of agents $M \subseteq N$ where $|M| = m \geq 2$, a circle of players C_m^t forms in M at time t if for any $i \in M$ there is a unique $j \in M$ such that $a_j^t = i$.*

In other words, a circle of players C_m^t is such that at time t all agents who belong to the subset M steal from each other.

In the following, we investigate the stealing game under various specifications for the parameters. We adopt a “bottom-up” approach, that is, we start from the simplest possible setting and progressively generalize the analysis. Our primary goal is to solve for the optimal timing decisions of the players. The fact that the stealing game is a dynamic game that may involve many players competing over many periods often renders unfeasible a complete characterization of the equilibria in terms of strategy profiles that define a complete plan of action for every possible contingency that may arise. Therefore, we will often define the equilibria of the game in terms of the action profiles that emerge on the equilibrium path rather than in terms of complete strategy profiles. In other words, we do not distinguish between strategy profiles that may differ in terms of off-path equilibrium behavior but still lead to the same equilibrium outcome.

4 The game with $T = 2$ and $K = 1$

We start the analysis of the game by focusing on a basic, yet highly informative, case. More precisely, we study a stealing game that lasts two periods (i.e., $T = 2$) and in which players have only one stealing opportunity (i.e., $K = 1$). We first analytically solve the game with two and three players, then extend the results to the case in which $n > 3$.

4.1 The two-player game

If $n = 2$, each player has only one opponent from which he can steal. The game is essentially a 2x2 game where players must only decide when to use their single stealing opportunity. The following proposition states that in the unique equilibrium of the game, both agents steal from each other in $t = 1$. The proof is trivial (in expectations stealing in $t = 1$ is strictly dominant) and is therefore omitted.

Proposition 4 *The stealing game with two players has a unique MPE. In this equilibrium, both agents belong to the circle C_2^1 .*

4.2 The three-player game

When $n = 3$, the stealing game presents multiple equilibria. These can be characterized as follows:

Proposition 5 *The three-player stealing game has four pure strategy MPE:*

- *two equilibria are such that every agent belongs to a circle C_3^1 (all players move in $t = 1$);*
- *two equilibria are such that every agent belongs to a circle C_3^2 (all players move in $t = 2$).*

Defining the set of players as $N = \{a, b, c\}$, Proposition 2 thus identifies the following equilibrium outcomes:

$$\begin{aligned} O_1 &= ((b, \emptyset), (c, \emptyset), (a, \emptyset)) & O_3 &= ((\emptyset, b), (\emptyset, c), (\emptyset, a)) \\ O_2 &= ((c, \emptyset), (a, \emptyset), (b, \emptyset)) & O_4 &= ((\emptyset, c), (\emptyset, a), (\emptyset, b)) \end{aligned}$$

All the equilibria are Pareto equivalent with $\bar{u}_i(\hat{\sigma}) = \frac{1}{3}$ for all i since all players are equally likely to finish the game as the largest shareholder. The interesting

feature of the three-player game, versus the two-player game, is that equilibria exist in which all agents postpone their moves. While in the two-player case, any strategy that prescribes a player to postpone his move is dominated, a similar relationship does not hold when the number of players equals three. There exist in fact states of the world where postponing one's move pays off. Consider, for instance, a situation in which player a is the unique agent who moves in the first period, and assume he steals from player b . Clearly, this is an ideal scenario for player c , because he can now observe the new state θ^2 and then decide how to use his stealing opportunity (which is still fully effective).

4.3 The game with $n > 3$ players

We now extend the analysis of the stealing game to a situation in which more than three players compete over two periods and have a single stealing opportunity. As before, our primary interest lies in investigating the timing of agents' moves and their decision whether to preempt or postpone their stealing opportunity. The following proposition defines the preempting equilibria, the ones in which all players move in $t = 1$.

Proposition 6 *The stealing game with more than three players has multiple pure strategy MPE in which all players move in $t = 1$. All these equilibria are such that every agent belongs to a circle of players.*

Notice that the number of preempting equilibria rapidly explodes with the number of players. In fact, for any $n > 3$, equilibrium profiles are not only those that support the $(n - 1)!$ possible circles that involve all the players (i.e., the circles C_n^1) but also those in which the set N is partitioned and smaller circles (possibly of different sizes) emerge in every part.

There also exist postponing equilibria in which all the players use their stealing opportunities in $t = 2$. However, the conditions that define them are stricter, as shown by the following proposition.

Proposition 7 *The stealing game with more than three players has multiple pure strategy MPE in which all players move in $t = 2$. All these equilibria are such that every agent belongs to a circle of players, and every circle contains at least three players.*

Proposition 4 indicates that there cannot exist postponing equilibria that feature circles made of two players because within any circle of this kind, players would like to deviate in order to preempt their rival. In fact, Proposition 1 showed that the only circle that qualifies as an equilibrium when $n = 2$ is the one in which both players move in $t = 1$.

Comparing Propositions 1 through 4, it is possible to state three additional results that characterize a stealing game in which $n \geq 2$ agents compete over two periods and have a single stealing opportunity:

- whenever $n \neq 3$, not all the strategy profiles where every agent belongs to a circle of players are equilibria.
- for any $n \geq 4$, the number of preempting equilibria is strictly larger than the number of postponing equilibria.⁴

⁴Consider, for instance, a stealing game with $n = 5$. Proposition 3 implies that preempting equilibria can emerge only in partitions (5) and (3, 2). The number of preempting equilibria is thus 44: there exist $4! = 24$ equilibrium outcomes in partition (5) and 20 equilibrium outcomes in partition (3, 2) (ten couples can be drawn from a set of 5 elements; for any of these couples there are two possible circles that can emerge in the part that involves 3 players). On the contrary, Proposition 4 states that postponing equilibria can emerge only in partition (5) since players must necessarily belong to a circle C_5^2 . It follows that there are only 24 postponing equilibria.

- for any $n \geq 4$, there exist equilibria that are asymmetric with respect to the timing decision.⁵

5 The game with $T > 2$ and $K < T$

As a further generalization, we consider a stealing game that lasts for T periods and where players have $K \in \{1, \dots, T - 1\}$ stealing possibilities. The widening of players' action spaces, paired with the stochastic nature of the game, rapidly enlarges the state space. This fact renders unfeasible a clear categorization of agents' best responses. As such, not only a proper characterization of the equilibria, but also the mere description of the action profiles that emerge along the various equilibrium paths, appear to be out of reach.

It is, however, possible to state some very general results. These maintain the same qualitative features as those presented in the previous sections, at least for what concerns the timing of agents' first move. The main insights are that in equilibrium, all players may remain idle for some initial periods (the two-player case being an exception), and that when a player uses his first stealing opportunity, he necessarily belongs to a circle of players. The following proposition formalizes these results:

Proposition 8 *All pure strategy MPE of a stealing game in which $n \geq 2$ players compete over $T > 2$ periods and have $K < T$ stealing opportunities are such that:*

- *if $n = 2$, both agents belong to the circles C_2^t for any $t = \{1, \dots, K\}$;*
- *if $n \geq 3$, each agent belongs to a circle $C_m^{t^*}$ where t^* is the period in which the agent uses his first stealing opportunity. In particular, $t^* = 1$ if $m = 2$ whereas $t^* \in \{1, \dots, T - K + 1\}$ if $m \geq 3$.*

⁵Let $N = \{a, b, c, d, e\}$. The outcome $O_1 = ((b, \emptyset), (a, \emptyset), (\emptyset, d), (\emptyset, e)(\emptyset, c))$ is an example of an asymmetric equilibrium: a and b belong to the circle C_2^1 and move in $t = 1$ while c , d , and e belong to a circle C_3^2 and move in $t = 2$.

6 Conclusions

We analyzed what we called the stealing game, a stochastic game in which players must decide when to steal portions of a homogeneous good from each other with the goal of finishing with the largest share. The peculiarity of the game is that the expected effectiveness of a player's theft is increasing in the agent's holdings. We showed that in a stealing game with two agents, players always want to preempt their rival and thus employ their stealing opportunities as soon as possible. Alternatively, we showed that with three players, the game also displays equilibria in which all the agents postpone their moves. Finally, we showed that when the number of players is larger than three, asymmetric equilibria exist, and not all the players necessarily steal in the same periods.

Acknowledgements

I thank an anonymous referee for very helpful comments that substantially improved the paper. I am also grateful to Pascal Courty, Dino Gerardi, Edoardo Grillo, Dorothea Kubler, Vilen Lipatov, Marco Mariotti, Ignacio Monzon, and Karl Schlag for useful suggestions and discussions. All remaining errors are my own.

Appendix

Proof of Proposition 2. We prove that the four outcomes are equilibrium outcomes by showing that no player has any incentive to deviate. Define the set of players as $N = \{a, b, c\}$ and denote the four candidate equilibrium outcomes as:

$$\begin{aligned}
O_1 &= ((b, \emptyset), (c, \emptyset), (a, \emptyset)) & O_3 &= ((\emptyset, b), (\emptyset, c), (\emptyset, a)) \\
O_2 &= ((c, \emptyset), (a, \emptyset), (b, \emptyset)) & O_4 &= ((\emptyset, c), (\emptyset, a), (\emptyset, b))
\end{aligned}$$

In what follows, we evaluate player a 's possible deviations from the strategies that support outcomes O_1 and O_3 . This implies no loss of generality since agents are initially symmetric and what we show about strategies that support O_1 and O_3 also holds for the strategies that support O_2 and O_4 . We indicate with $\hat{\sigma}$ a strategy profile that leads to outcome O_x with $x \in \{1, 3\}$. The resulting expected allocation $\bar{\pi}^3(\hat{\sigma}) = (\bar{\pi}_a^3(\hat{\sigma}), \bar{\pi}_b^3(\hat{\sigma}), \bar{\pi}_c^3(\hat{\sigma}))$ is given by

$$\bar{\pi}^3(\hat{\sigma}) = \left(\frac{1}{3} + \bar{y}_a^t - \bar{y}_c^t, \frac{1}{3} + \bar{y}_b^t - \bar{y}_a^t, \frac{1}{3} + \bar{y}_c^t - \bar{y}_b^t \right) \text{ with } t \in \{1, 2\} \quad (3)$$

Because of Assumption 1, Lemma 2, and the fact that agents are symmetric and move simultaneously, it follows that $\bar{y}_i^t = \bar{y}_j^t$ for any $i, j \in \{a, b, c\}$. Therefore, $\bar{\pi}_i^3(\hat{\sigma}) = \bar{\pi}_j^3(\hat{\sigma})$ for any $i, j \in \{a, b, c\}$ and the expected payoff of agent a is $\bar{u}_a(\hat{\sigma}) = \frac{1}{3}$.

By Lemma 1, all deviations such that player a does not use his stealing opportunity are dominated. We thus only need to show that no deviations in which a uses his stealing opportunity are profitable. Agent a can deviate from $\hat{\sigma}_a$ by stealing from a different rival (we analyze this possibility in Case 1 below) and/or by stealing in a different period (Cases 2a and 2b).

Case 1) Consider any outcome O_x , $x \in \{1, 3\}$, and let player a deviate and steal from c rather than from b . We denote such a deviation by $\tilde{\sigma}_a$ and the resulting profile by $\tilde{\sigma} = (\tilde{\sigma}_a, \hat{\sigma}_{-a})$. The final expected allocation is given by

$$\bar{\pi}^3(\tilde{\sigma}) = \left(\frac{1}{3} + \bar{y}_a^t - \bar{y}_c^t, \frac{1}{3} + \bar{y}_b^t, \frac{1}{3} + \bar{y}_c^t - \bar{y}_a^t - \bar{y}_b^t \right) \text{ with } t \in \{1, 2\} \quad (4)$$

Notice that in $\tilde{\sigma}$ all players still move simultaneously and thus, by Assumption 1, $\bar{s}_i^t = \bar{s}_j^t$ for any $i, j \in \{a, b, c\}$. By expression (1) and Lemma 2, it then follows

that $\bar{y}_a^t = \bar{y}_b^t$ whereas $\bar{y}_a^t \leq \bar{y}_c^t$ (strict inequality holds when $\bar{s}_i^t > \frac{1}{6}$, in which case $\bar{y}_a^t = \bar{y}_b^t = \frac{1}{6}$ as agents a and b share c 's initial endowment). Therefore, $\bar{\pi}_a^3(\tilde{\sigma}) \leq \frac{1}{3}$. Consider first the case $\bar{\pi}_a^3(\tilde{\sigma}) = \frac{1}{3}$ and observe that $\bar{\pi}_b^3(\tilde{\sigma}) > \frac{1}{3}$. Given that $\bar{\pi}_a^3(\tilde{\sigma}) = \bar{\pi}_a^3(\hat{\sigma})$ and $\bar{\pi}_b^3(\tilde{\sigma}) + \bar{\pi}_c^3(\tilde{\sigma}) = \bar{\pi}_b^3(\hat{\sigma}) + \bar{\pi}_c^3(\hat{\sigma}) = \frac{2}{3}$, it follows that player a 's probability of winning the game (i.e., the probability of $\pi_a^3(\cdot)$ being larger than *both* $\pi_b^3(\cdot)$ and $\pi_c^3(\cdot)$) is smaller with $\tilde{\sigma}$ rather than with $\hat{\sigma}$. It follows that $\bar{u}_a(\tilde{\sigma}) < \frac{1}{3}$. The same result holds a fortiori if $\bar{\pi}_a^3(\tilde{\sigma}) < \frac{1}{3}$. Therefore, the deviation to $\tilde{\sigma}_a$ is unprofitable.

Case 2a) Consider outcome O_1 and let player a evaluate the possibility to postpone his stealing attempt to the second period. We denote such a deviation by $\check{\sigma}_a$ and the resulting profile by $\check{\sigma} = (\check{\sigma}_a, \hat{\sigma}_{-a})$. The expected allocation at the beginning of $t = 2$ is given by

$$\bar{\pi}^2(\check{\sigma}) = \left(\frac{1}{3} - \bar{y}_c^1, \frac{1}{3} + \bar{y}_b^1, \frac{1}{3} + \bar{y}_c^1 - \bar{y}_b^1 \right) \quad (5)$$

with $\bar{y}_b^1 = \bar{y}_c^1$ and $\bar{\pi}_a^2(\check{\sigma}) \geq 0$. In expectations, agent a will thus target player b as $\bar{\pi}_b^2(\check{\sigma}) > \bar{\pi}_c^2(\check{\sigma})$. If $\bar{\pi}_a^2(\check{\sigma}) = 0$ then, by Assumption 1, $\bar{s}_a^2 = 0$. Therefore, $\bar{y}_a^2 = y_a^2 = 0$, $\bar{\pi}_a^3(\check{\sigma}) = 0$ and $u_a(\check{\sigma}) = 0$. If instead $\bar{\pi}_a^2(\check{\sigma}) > 0$ then $\bar{y}_a^2 > 0$ and the expected final allocation will be

$$\bar{\pi}^3(\check{\sigma}) = \left(\frac{1}{3} - \bar{y}_c^1 + \bar{y}_a^2, \frac{1}{3} + \bar{y}_b^1 - \bar{y}_a^2, \frac{1}{3} + \bar{y}_c^1 - \bar{y}_b^1 \right) \quad (6)$$

with $\bar{y}_a^2 < \bar{y}_i^1$ for any $i \in \{b, c\}$ such that $\bar{\pi}_a^3(\check{\sigma}) < \bar{\pi}_c^3(\check{\sigma}) < \bar{\pi}_b^3(\check{\sigma})$. Therefore, $\bar{u}_a(\check{\sigma}) < \frac{1}{3}$ and the deviation is unprofitable.

Case 2b) Consider outcome O_2 and let player a deviate and steals from b in $t = 1$ rather than in $t = 2$ (the same reasoning applies in case a decides to steal from c). We denote such a deviation by $\ddot{\sigma}_a$ and the resulting profile by $\ddot{\sigma} = (\ddot{\sigma}_a, \hat{\sigma}_{-a})$. The expected allocation at the beginning of $t = 2$ is given by

$$\bar{\pi}^2(\ddot{\sigma}) = \left(\frac{1}{3} + \bar{y}_a^1, \frac{1}{3} - \bar{y}_a^1, \frac{1}{3} \right) \quad (7)$$

In $t = 2$ player c will thus steal from a . Player b with $\bar{\pi}_b^2(\ddot{\sigma}) \geq 0$ takes this into account and decides whether to steal from a or c , depending on the actual realization y_a^1 . If b steals $\bar{y}_b^2 \geq 0$ from c , then the expected final allocation is

$$\bar{\pi}^3(\ddot{\sigma}) = \left(\frac{1}{3} + \bar{y}_a^1 - \bar{y}_c^2, \frac{1}{3} - \bar{y}_a^1 + \bar{y}_b^2, \frac{1}{3} + \bar{y}_c^2 - \bar{y}_b^2 \right) \quad (8)$$

with $\bar{y}_a^1 = \bar{y}_c^2$ because of Assumption 1 and Lemma 2. Therefore, $\bar{\pi}_b^3(\ddot{\sigma}) < \bar{\pi}_a^3(\ddot{\sigma}) < \bar{\pi}_c^3(\ddot{\sigma})$ with $\bar{\pi}_a^3(\ddot{\sigma}) = \frac{1}{3}$. By the same reasons discussed in Case 1, it follows that $\bar{u}_a(\ddot{\sigma}) < \frac{1}{3}$. If instead, b steals from a , then the expected final allocation is

$$\bar{\pi}^3(\ddot{\sigma}) = \left(\frac{1}{3} + \bar{y}_a^1 - \bar{y}_c^2 - \bar{y}_b^2, \frac{1}{3} - \bar{y}_a^1 + \bar{y}_b^2, \frac{1}{3} + \bar{y}_c^2 \right) \quad (9)$$

and, a fortiori, $\bar{u}_a(\ddot{\sigma}) < \frac{1}{3}$. Therefore, the deviation to $\ddot{\sigma}_a$ is unprofitable.

In summary, all possible deviations are strictly unprofitable and we can thus conclude that the four outcomes are indeed equilibrium outcomes. Moreover, in discussing the possible deviations of generic player a , we have spanned all possible outcomes of the game. In addition to outcomes O_x , $x \in \{1, 2, 3, 4\}$, all remaining outcomes belong in fact to at least one of these categories: i) at least one player does not use his stealing opportunity (these outcomes cannot be equilibria outcomes because of Lemma 1); ii) all players move in the same period but do not belong to a circle C_3^t with $t \in \{1, 2\}$ (these outcomes cannot be equilibria outcomes as discussed in Case 1); iii) two players move in period $t \in \{1, 2\}$ while the third agent moves in period $t' \neq t$ (these outcomes cannot be equilibria outcomes as discussed in Cases 2a and 2b). It follows that the outcomes O_x , $x \in \{1, 2, 3, 4\}$, are the unique equilibrium outcomes. Therefore, the strategy profiles that support these outcomes are the unique pure strategy Markov perfect equilibria of the game. ■

Proof of Proposition 3. The proof generalizes parts of the proof of Proposition 2 (in particular, Cases 1 and 2a). Take any strategy profile $\hat{\sigma}$ that supports an outcome in which each agent uses his stealing opportunity in $t = 1$ and belongs to a circle of players. Notice that $\bar{u}_i(\hat{\sigma}) = \frac{1}{n}$ for any $i \in N$. Say that, according to $\hat{\sigma}$, generic player a belongs to a circle C_m^1 where he steals from agent b . Now let a evaluate possible deviations. If a moves in $t = 1$ but robs agent $c \neq b$ (it does not matter whether c also belongs to C_m^1 or not) then, in the modified profile $\tilde{\sigma} = (\tilde{\sigma}_a, \hat{\sigma}_{-a})$, player b is not robbed by anyone. The expected final allocation $\bar{\pi}^3(\tilde{\sigma})$ is such that $\bar{\pi}_a^3(\tilde{\sigma}) = \bar{\pi}_a^3(\hat{\sigma}) = \frac{1}{n} + \bar{y}_a^1 - \bar{y}_d^1$ (d is the player who robs a in $\hat{\sigma}$ and in $\tilde{\sigma}$). However, $\bar{\pi}_b^3(\tilde{\sigma}) = \frac{1}{n} + \bar{y}_b^1$ such that $\bar{\pi}_a^3(\tilde{\sigma}) < \bar{\pi}_b^3(\tilde{\sigma})$. It follows that $\bar{u}_a(\tilde{\sigma}) < \frac{1}{n}$ and the deviation to $\tilde{\sigma}_a$ is unprofitable. If instead player a postpones his move and steals from b (or any other player) in $t = 2$ (we indicate such a strategy with $\check{\sigma}_a$), then $\bar{\pi}_a^3(\check{\sigma}) = \frac{1}{n} - \bar{y}_d^1 + \bar{y}_a^2$ such that, by Assumption 1 and Lemma 2, $\bar{\pi}_a^3(\check{\sigma}) < \frac{1}{n}$. Therefore, $\bar{u}_a(\check{\sigma}) < \frac{1}{n}$ and the deviation to $\check{\sigma}_a$ is unprofitable. ■

Proof of Proposition 4. The proof generalizes parts of the proof of Proposition 2 (in particular, Cases 1 and 2b). Consider any strategy profile $\hat{\sigma}$ that supports an outcome in which each agent uses his stealing opportunity in $t = 2$ and belongs to a circle of players C_m^2 with $m \geq 3$. Agent a 's possible deviations (a steals from a different player in $t = 2$, or brings his move forward to $t = 1$) are unprofitable (see the proofs of Propositions 2 and 3). As for the restriction on the minimum size of any circle of players, consider a circle with two players C_2^2 . The strategic situation in C_2^2 replicates the one that characterizes the stealing game when $n = 2$. As such (see Proposition 1), no circle C_2^2 can be part of any equilibrium as both players would like to deviate and use their stealing opportunity in $t = 1$. ■

Proof of Proposition 5. The proof is trivial for the case with $n = 2$: both players use all their stealing attempts as soon as possible (and thus belong to

circles C_2^t for any $t = \{1, \dots, K\}$) as this is a dominant strategy. When $n \geq 3$, a necessary condition for any strategy profile to qualify as an equilibrium profile is that each agent belongs to a circle of players at the moment he moves for the first time. Indeed, if this was not the case, there would certainly exist at least one player who could profitably deviate (see the proofs of Propositions 2, 3, and 4). Moreover, agents that belong to circles of two players must necessarily use their first stealing opportunity in period $t' = 1$ (see Proposition 1). Agents that belong to circles of more than two players can instead use their first stealing opportunity at any $t' \in \{1, \dots, T - K + 1\}$ where the upper bound is due to the fact that, by Lemma 1, players plan to use all their stealing attempts throughout the game.

■

References

- [1] Borel, E., 1921. La theorie du jeu les equations integrales a noyau symetrique. Comptes Rendus de l'Academie 173, 1304–1308. English translation by Savage, L., 1953. The Theory of Play and Integral Equations with Skew Symmetric Kernels. Econometrica 21, 97–100.
- [2] Brunnermeier, M.K., Morgan, J., 2010. Clock games: theory and experiments. Games and Economic Behavior 68, 532-550.
- [3] Bulow, J., Klemperer, P., 1999. The generalized war of attrition. American Economic Review 89, 175-189.
- [4] Dubovik, A., Parakhonyak, A., 2014. Drugs, guns, and targeted competition. Games and Economic Behavior 87, 497-507.
- [5] Harsanyi, J., Selten, R., 1988. A general theory of equilibrium selection in games. MIT Press, Cambridge.

- [6] Konrad, K.A., 2009. Strategy and dynamics in contests. Oxford University Press, Oxford, UK.
- [7] Maynard Smith, J., 1974. Theory of games and the evolution of animal contests. *Journal of Theoretical Biology* 47, 209-221.
- [8] Maskin, E., Tirole, J., 2001. Markov perfect equilibrium I. Observable actions. *Journal of Economic Theory* 100, 191-219.
- [9] Park, A., Smith, L., 2008. Caller number five and related timing games. *Theoretical Economics* 3, 231-256.
- [10] Rinott, Y., Scarsini, M., Yu, Y., 2012. A Colonel Blotto gladiator game. *Mathematics of Operations Research* 37, 574-590.
- [11] Rosenthal, R., 1981. Games of perfect information, predatory pricing, and the chain store paradox. *Journal of Economic Theory* 25, 92-100.
- [12] Sela, A., Erez, E., 2013. Dynamic contests with resource constraints. *Social Choice and Welfare* 41, 863-882.
- [13] von Stackelberg, H., 1934. *Marktform und Gleichgewicht*. Vienna and Berlin: Springer Verlag.