Split it up to Create Incentives: Investment, Public Goods and Crossing the River

Simon Martin[*] Karl H. Schlag[◊]

March 27, 2019

Abstract

Many allocations cannot be implemented as each of the parties involved has an incentive to deviate from what they need to contribute. Leading examples are sequential trade, provision of a public good or when to pay the ferryman when crossing river Styx. We show how to implement a desired allocation by breaking contributions into pieces, each party threatening to discontinue if others deviate. Contributions are typically designed to decrease over time. In the case of river Styx, the ferryman needs to be paid right before arrival. Our solution concept is ε-subgame perfect equilibrium where players only deviate when alternative choices lead to substantially higher payoffs.

Keywords: holdup problem, ε-subgame perfect equilibrium, finite horizon, enforcement without contracts, gradualism

JEL Codes: D23, C72, L14

[*]University of Vienna. Email: simon.martin@univie.ac.at

[◊]University of Vienna. Email: karl.schlag@univie.ac.at

1 Introduction

"Don't pay the ferryman,
Don't even fix a price,
Don't pay the ferryman,
Until he gets you to the other side"

Chris de Burgh, 1982

Chris de Burgh's 1982 pop song 'Don't pay the ferryman' is a reference to the ferryman Charon from Greek mythology, who took the deceased from one side of the river Styx to the other side, in exchange for a small fee (Nardo, 2002). When should one pay the ferryman? According to Chris de Burgh, not until he gets you to the other side, since he doesn't have an incentive to continue the passage once you have paid. However, payment at the other side cannot be the solution to the problem as the ferryman knows that the passenger has no incentive to pay once he has reached the other side.

There are many real world settings that exhibit the same flavor. When should we pay for the delivery of a good when trade is sequential? If payment occurs after delivery then there are no incentives to pay which deprives the seller of the incentive to deliver in the first place. With payment prior to delivery there is no incentive to deliver. Similarly, how can we convince players to contribute to a public good when it is a dominant strategy to abstain from contribution?

The common feature of these examples that can be found in numerous economic applications is that investments are irreversible and desirable allocations cannot be implemented in equilibrium without additional commitment power. In particular, in our above examples the subgame perfect equilibrium prediction is that the passenger will not be transported, trade will not take place and the public good will not be provided.

There are nevertheless ways to implement desirable allocations against individual incentives to deviate. Most prominently, institutions may be used to design enforceable contracts to ensure provision. In addition, repeated interaction and reputational concerns may create sufficiently strong incentives to build up a lasting relationship[1]. Moreover, as shown by Pitchford and Snyder (2004), if the duration of the relationship between a buyer

[1]For an early theoretical contribution see Shapiro (1982), and Dulleck et al. (2011) and Palfrey and Prisbrey (1996) for more recent experiments.

and seller is uncertain and potentially never ending one can overcome incentive problems by splitting the entire investment into smaller pieces. However, this last solution is less practical as one has to split the transaction into infinitely many pieces and the entire good is never completely delivered. Note that the first two solutions also have substantial drawbacks. They limit the ability to trade in modern market economies defined by globalization and decentralization, as institutions are costly and often not available and potential trading partners are constantly changing.

In this paper, we propose a novel formal model for how to solve these allocation problems. Without institutions and within a single encounter we can get the passenger across the river, ensure trade at the agreed price and get everyone to contribute to the public good. Our approach gives insights into existing practices and qualifies how to better design interactions. Most importantly, our method is particularly simple to implement. We solve the problem by splitting up the total contribution (e.g. investment or payment) into finitely many smaller parts that typically decrease in size. In many applications only two rounds of contribution are needed. Non-compliance is punished by the termination of the relationship. The resulting strategies implement the desirable allocation in a variant of a subgame perfect equilibrium in which players only deviate if this leads to substantial gains. The formal solution concept is that of an ε-subgame perfect equilibrium (Mailath et al., 2005).

In a subgame perfect equilibrium (SPE) no player can unilaterally obtain strictly higher payoffs in any subgame by deviating. The concept of an ε-SPE makes a minor adjustment. There is a constant $\varepsilon > 0$ which is typically small, such that no player can obtain at least ε more in any subgame by deviating. Intuitively, why bother about arbitrarily small gains, figuratively often referred to as peanuts? The ε threshold can capture deliberation costs and costs of embarrassment costs due to deviating from the suggested plan.

The method for how to cut up the total contribution is simple. To avoid unravelling from the end one starts the construction with the final contribution. One looks for some final contribution that involves a low cost for the party who makes it and involves high benefits for the others. Ideally the benefits to others are substantial. For instance, this is the case when handing over the key to a house or a car, writing a reference letter or opening the door to let the passenger off the boat. The last mover is believed to deliver as this only involves a small cost for her. Whether or not the last mover delivers means a lot to others who hence are willing to incur large costs in return earlier in the game. If the benefits of the last mover to others are sufficiently large then one only has to split the entire contribution into two. For instance, in the case of the river passage, let the last contribution be the ferryman letting the passenger off the boat at the dock. This is a small cost to the ferryman and generates large benefits to the passenger. The passenger anticipates this last move and is willing to pay the ferryman the entire fee upon arrival at the dock before getting off the boat. Consequently, the ferryman does not ask for an upfront payment. Both sides receive their preferred outcomes. If benefits to others of the final contribution are not that large then one has to split up the total contribution into more parts. The size of the contributions of each player typically decreases over time as early on each player has more to lose by deviating. In particular, in the context of this model, it does not make sense to split the total contribution into equally sized pieces, as is common practice in job order contracting. If the pieces are not small, the contribution of the last mover will not be credible. If the pieces are small then there is excessive splitting as all pieces have the same size.

We first demonstrate the basic insights in the simplest setting with arbitrary actions and one or two choices per player. Then we present the general construction for the case where the choice (also called investment) of each player has a single dimension (e.g., price, observable effort). Payoffs are differentiable and only a function of the aggregate investments of each player. In this setting, individual investments are irreversible and assumed to be costly for the player investing but beneficial for the others. Players move alternatingly and only care about their final payoff (in a later section we discuss the case of discounting). First we look at the case of two players. The objective is to reach a given target as defined by a given level of aggregate investments for each player. This is done by specifying a sequence of investments such that if any player deviates then both players discontinue to invest. The idea is to specify investments in a way that can be supported by an ε-SPE. The (incentive compatability) constraints that ensure that the desired investments are made are determined as follows. Each player looks before each investment at the total benefit of all investments and compares this to the value of investments made in the past.

An intuitive way to construct the sequence of investments is to make all incentive compatability constraints after the first round binding. Under an additional condition on payoffs this construction can be used to reach any target and to do this in a minimal number of rounds. Intuitively, we need that larger later investments enable larger earlier investments and thus reduce the total number of rounds needed. This condition is ensured if the payoff of a given player increases when all players marginally increase their investment. The investment levels in this construction are determined by a set of recursive equations. These are easy to explicitly solve in the buyer-seller setting. In this model, one player as seller makes a costly investment that benefits the other player as buyer who pays the seller for this investment. Our solution recommends for the buyer at the start to make an upfront payment to the seller, thereafter to reimburse the seller for how much he just invested and to make a last payment equal to ε. We refer to this as a pay-as-you-go system. The behaviour of the seller is similar. In each round the seller creates an added value to the buyer that is exactly equal to her last payment. In the buyer-seller model we can also easily incorporate any degree of discounting.

We then move to more than two players and identify a sufficient condition for the above construction to reach any given target with the minimal number of investments. We also present analogous results for simultaneous move games among symmetric players. In a separate section we show in numerical examples how the number of investments depends on the threshold ε and on the payoff functions. In a further section we demonstrate how our insights extend to the finitely repeated Prisoners' Dilemma if it is possible to change the frequency of interaction between the players. In particular, we show how grim trigger strategies can incentivize cooperation if the time between interactions becomes shorter towards the end of game. Finally we show how social preferences, introduced instead of the ε threshold, suffice to implement desirable investments in finite horizons.

Splitting up an investment was advocated in a different context by Dixit and Nalebuff (1993) to better deal with incomplete information. In their case, splitting serves as a tool to limit possible future losses if the business partner turns out not to be trustworthy. There is no suggestion on how to split up the investment, in particular there is no formal model in that paper.

There is a series of papers that investigate the choice of irreversible actions over time. In Admati and Perry (1991), agents care for completion of a project within a finite time horizon and hence the final investments can be self enforcing. In contrast, investments are costly in our setting. The other papers incentivize costly actions by considering a never ending relationship with an infinite time horizon. Irreversible choices appear as investments in the infinitely repeated game modelled in Lockwood and Thomas (2002), as concessions in the bargaining model of Compte and Jehiel (2004), and as both investment and bargaining in Che and Sákovics (2004)

For other applications of ε-optimality see e.g. Radner et al. (1980), Baye and Morgan (2004), Barlo and Dalkiran (2009) and Milgrom (2010).

We proceed as follows. In Section 2 we informally present the basic principles using two players where at most one of them moves twice. In Section 3, we first present the main model for alternating investments. Results for two-player games are in Section 3.2 and results for games with more than three players are in Section 3.5. We extend our results to simultaneous move settings in Section 4 and show numerical examples in Section 5. Extensions for social preferences, discounting and costs of splitting and the Prisoners' Dilemma are shown in Section 6. In Section 7 we conclude.

2 An informal view

In this section we illustrate the problem and how we solve it, using minimal formal notation and starting from the simplest case. In the main part of the paper we introduce the general model.

Consider a set of players who move alternatingly in a given order. When a player moves she can either choose to be inactive or she can choose an action that is costly to herself but beneficial to others. This action might be an investment in a relationship, a contribution to a public good or the delivery of one of the parts of a good. In the following we will refer to actions as investments.

There is complete and perfect information. So each player observes all past investments and knows the payoffs of each other player. Each time a player moves she cannot revoke her earlier investments.

Consider some target that consists of a set of investments to be made by each player such that if all players make these then each player is better off than if all players remain inactive. The objective is to try to find a sequence of investments that yields the desired target, and that incentivizes each player in the end to choose the investments associated to the target. This sequence of investments will be implemented by each player stopping to invest if some player in the past did not do their agreed share, so did not make choices as prescribed by this sequence.

The problem is that whoever moves last will not wish to make that final investment in a subgame perfect equilibrium (SPE) as investments are costly to the person investing. Benefit is only for others. In this paper we consider players that will follow the suggested strategy even if other choices are better, as long as these do not generate substantially higher payoffs. So we will employ the solution concept of ε-SPE. Hence if the last investment is not very costly to that player then she will make it. However we need to see if this can incentivize those moving earlier to also make their investments. To understand the mechanisms we first start with the case where there are only two players.

Assume that player one moves first, followed by player two. The objective is to incentivize player one and player two to choose a given pair of investments that makes both better off than when they both remain inactive. This is done as follows. Player two promises to make the desired investment if player one also does and threatens to be inactive otherwise. For the promise of player two to follow suit to be credible we need an extra condition, namely that, given the desired investment of player one, the cost of the desired investment of player two is negligible. Note that the threat of player two to otherwise remain inactive is credible. As investments are costly, it is always a best response for player two to remain inactive. For player two to remain inactive means that player one is worse off than when being inactive. Hence, this combination of promise and threat induces player one to choose the desired investment.

An example where these conditions hold is the following. Player one as worker is first supposed to make a given observable effort that benefits player two as employer. Then player two is supposed to write a positive reference letter for player one. Similarly, the owner of an apartment is supposed to make everything satisfactory to the guest who is supposed to send the owner a refrigerator magnet once arriving back home[2]. In both cases we see that the addition of the reference letter or of the magnet helps to incentivize player one to make the desired effort. An alternative example is the ferry passage. It works if the price of the passage is negligible from the perspective of the passenger. In this case,

[2]Note that asking for a magnet can yield stronger incentives than only asking for a positive online review as each only has a marginal effect in case the owner already has many reviews.

the ferryman moves first and decides whether or not to take the passenger across the river, followed by the passenger who decides whether or not to pay the agreed price. It also works if the effort of the ferryman to bring the passenger across is negligible. In that case, the passenger first pays the fee and then the ferryman takes him or her across.

For completeness, here are formal conditions. Let $\varepsilon > 0$ be the threshold under which profits of deviation are considered negligible. Let A_i be the set of actions of player i, with $A = A_1 \times A_2$, where A_i contains the possibility of inaction, denoted by n_i. Let $u_i(a)$ be the utility of player i given the profile of actions a for $i = 1, 2$. Investments are costly, so $u_1(n_1, a_2) \geq u_1(a)$ and $u_2(a_1, n_2) \geq u_2(a)$ for $a \in A$ and $i = 1, 2$. Let $\bar{a} \in A$ be the desired choice of actions where $u_i(\bar{a}) \geq u_i(n)$ for $i = 1, 2$. Then the sufficient condition for implementing $\bar{a}$ is that $u_2(\bar{a}_1, \bar{a}_2) \geq u_2(\bar{a}_1, n_2) - \varepsilon$.

Now consider the case where the desired investments bear substantial cost for each player. For instance, assume that the employer also has to pay the worker a wage. Or consider the case where both the effort of the ferryman and the fee of crossing the river are not negligible. Then the above construction does not work. The last mover will not invest and consequently neither will the first. However we might be able to incentivize the desired investments if the choice of player two can be split up into two parts. The idea is to let player two move first, followed by player one and then by player two moving again. As before, to incentivize the last move of player two, the desired second investment of player two has to involve negligible costs when there was no previous deviation. To incentivize the investment of player one we need that the total benefits for player one exceed the benefits to her when only player two makes the desired investment in the first round and all other choices are inaction. Note that the first investment of player two is incentivized as deviation is punished by player one choosing inaction.

For example, the investments of player one as worker can be to work hard, of player two as employer to pay a good wage and to write a good reference letter for player one. The employer pays the good wage as she benefits from the effort of the worker. The employer writes the reference letter as the time she invests in writing the letter is really not a big deal. The worker makes the effort if she prefers to work and get a good reference letter to not working but running off with the wage and forfeiting the good reference letter. In case of the river crossing we have the following scenario. The ferryman is player two who first takes the passenger across the river. Then the passenger as player one pays the entire fee, and finally, the ferryman lets the passenger off the boat. Note that the ferryman example works regardless of how difficult it is to cross the river, as long as the effort for the ferryman of letting the passenger off the boat is negligible. Some real life examples that have a similar flavor to the ferryman example include ordering a good and paying on delivery before it is unloaded from the truck, and buying a car by first receiving the official papers, then transferring the money and finally getting the key to the car.

A bit more formally, we maintain the above notation, now being more specific about the actions of player 2, namely each action a_2 of player two consist of two parts, so $A_2 \subseteq B_1 \times B_2$ where B_k is the investment made at the kth choice of player two with B_k containing the possibility of inaction, $k = 1, 2$. Investments of player one and those of player two in the last round are costly, so $u_1(n_1, a_2) \geq u_1(a)$ and $u_2(a_1, (a_{21}, n_2)) \geq u_2(a)$ for all $a \in A$. Then a sufficient condition for implementing $\bar{a} \in A$ is that $u_2(\bar{a}) \geq u_2(\bar{a}_1, (\bar{a}_{21}, n_2)) - \varepsilon$, $u_1(\bar{a}) \geq u_1(n_1, (\bar{a}_{21}, n_2)) - \varepsilon$ and $u_2(\bar{a}) \geq u_2(n_1, (n_2, n_2))$.

It is of course possible that the first investment of player two is so large and beneficial to player one that player one is no longer willing to invest. In our example this may be the case if the initial wage together with the value of shirking is substantially larger than the value of the good reference letter after working hard. In that case we need more rounds of alternating investment. This is the subject of the remainder of the paper.

3 Alternating investments

3.1 Model

Consider a n-player game with alternating moves and perfect information. We denote the t-th round as the one which is t rounds from the end of the game, and the total number of rounds is denoted as T. Denote as player $i = \{1, 2, \ldots, n\}$ the player who is making the i-th move in each round. The order of moves in each round is fixed ex-ante. In each round t, each player chooses an action $x_{i,t} \in \mathbb{R}_+$. The sequence of investments (alternatively also referred to as *investement schedule*) across all rounds is thus written as $(x_t)_{t=1}^T$ where $x_t := (x_{i,t})_{i=1}^n$. Restricting consideration to positive actions reflects the idea of incrementally investing or contributing. We refer to either player's action as 'investments' in the following. Alternatively, this can be thought of as irrevocable levels of cooperation (Lockwood and Thomas, 2002). Each player can simply pick the outside option whenever he moves, which ends the game immediately. Otherwise, payoffs are accrued after the last move.

Defining as x^j the total investments of player j, payoffs for player i are given by $u_i\left((x^j)_{j=1}^n\right)$. Payoffs are continuous and differentiable and strictly decreasing in own investment, but increasing in the investments of the other players:

Assumption 1 (Payoffs and investments)**.**

$$\partial u_i/\partial x^i < 0 \text{ } \textit{for } i \neq j$$

$$\partial u_i/\partial x^j \geq 0 \text{ } \textit{for } i \neq j$$

We also assume that each player benefits when all investments increase equally:

Assumption 2 (Increasing utility by uniform investment increase)**.**

$$\sum_{j=1}^{n} \frac{\partial u_k}{\partial x^j} \geq 0 \ \forall k.$$

Notice that this implies that $u_k(y, y, \ldots, y)$ is increasing in y. We are interested in implementing a target $\bar{x} := (\bar{x}^j)_{j=1}^n$, i.e., an equilibrium such that the total investments of each player j equal $\bar{x}^j$. This target may be the socially optimal level of investment in certain applications, but this need not be the case. We are agnostic about the source of this target.

Our solution concept is ε-SPE (Mailath et al., 2005), which applies ε-Nash equilibrium to all subgames. A player only deviates if he gains more than a given threshold ε. Formally, the concept is defined as follows:

Definition 1. *Denoting as A_i the set of actions for player i and as σ^*_{-i} the strategy profile of the other players, a strategy profile σ^* is an **ε-Nash equilibrium** if $u_i(a_i, \sigma^*_{-i}) \leq u_i(\sigma^*) + \varepsilon$, $\forall a_i \in A_i$, $\forall i$.*

Definition 2. *A strategy profile σ^* is an **ε-subgame perfect equilibrium** if it induces an ε-Nash equilibrium in every subgame.*

Possible reasons why one might not want to deviate when the gains are small or negligible include: i) deliberation costs, ii) costs from being embarrassed (or alternatively, extra utility from staying at the proposed equilibrium) and iii) existence of possibility to prevent small deviations (e.g., small retaliation is available).

In the following we will limit attention to equilibria in which any deviation is punished by discontinuing the relationship. There is no loss of generality in this assumption when searching for outcomes that can be sustained in some ε-SPE. A necessary and sufficient condition for a sequence of investments to be supported in an ε-SPE is that neither player has an incentive to deviate in any round t. For player i, this requires that

$$u_i\left((\bar{x}^j)_{j=1}^n\right) \geq u_i\left(\left(\bar{x}^j - \sum_{k=1}^{t-1} x_{j,k} - x_{j,t}\mathbb{1}_{\{j\leq i\}}\right)_{j=1}^n\right) - \varepsilon \tag{1}$$

where $\mathbb{1}$ denotes the indicator function. The left-hand side represents the player's payoff when there is no deviation, and the right-hand side the player's payoff attained when choosing not to invest in round t.

3.2 Games with two players

Having presented the general framework, we first restrict attention to two-player games. This allows us to highlight the key mechanism and to show that our approach that involves making incentive compatibility constraints binding in all rounds allows to implement the target as fast as possible.

Given our model setup with ε-SPE and allowing the total transaction to be split across multiple rounds, implementing of the target is actually not that difficult. Indeed, an investment schedule with constant investments for each player in each round is guaranteed to eventually implement the target (see Proposition 6 in the appendix). The drawback is that this approach might take very long. For instance, for $u_1(x^1, x^2) = -x^1 + x^2$ and $u_2(x^1, x^2) = 2\sqrt{x^1} - x^2$, it takes 50 rounds to implement the target $(1, 1)$ when $\varepsilon = 0.01$, i.e., 1% of the potential total surplus.

The binding constraint is player two at the very end; there is nothing left to be gained for him by investing, so the maximal investment he is willing to make cannot make him more than ε worse off than not investing at all. But matters are different when it's his turn again, because than he is incentivized to invest himself by the prospect of player one also investing. A constant investment does not incorporate these dynamically changing incentive compatibility constraints.

The main result of this section is the following. Therein we consider an investment schedule in which incentive compatibility constraints are binding in each round. This schedule is guaranteed to implement the target within T rounds. Moreover, this schedule is optimal, in the sense that one cannot implement the target with less investment rounds.

Proposition 1. *Consider a sequential two-player game and fix $\varepsilon > 0$ and a target $\bar{x}$.*

i) There exists a T and an investment schedule $(x_t)_{t=1}^T$ such that incentive compatibility constraints (1) are binding for $t < T$ for each player and the target $\bar{x}$ is implemented in an ε-SPE.

ii) There does not exist a $T' < T$ such that $\bar{x}$ can be implemented in an ε-SPE that involves T' rounds.

Proof. i) By definition, neither player has an incentive to deviate in any round when the incentive compatibility constraints (1) is binding. We now show that whenever this is the case, investments are increasing over time. Let $x_{2,1}$ solve

$$u_2(\bar{x}^1, \bar{x}^2) = u_2(\bar{x}^1, \bar{x}^2 - x_{2,1}) - \varepsilon, \tag{2}$$

define $x_{1,1}$ as the solution to

$$u_1(\bar{x}^1, \bar{x}^2) = u_1(\bar{x}^1 - x_{1,1}, \bar{x}^2 - x_{2,1}) - \varepsilon \tag{3}$$

and finally define $x_{2,2}$ as the solution to

$$u_2(\bar{x}^1, \bar{x}^2) = u_2(\bar{x}^1 - x_{1,1}, \bar{x}^2 - x_{2,1} - x_{2,2}) - \varepsilon. \tag{4}$$

We first show that $x_{1,1} \geq x_{2,1}$. Starting from (3) we have that

$$\begin{aligned} -\varepsilon &= u_1(\bar{x}^1, \bar{x}^2) - u_1(\bar{x}^1 - x_{1,1}, \bar{x}^2 - x_{2,1}) \\ &< 0 \leq u_1(\bar{x}^1, \bar{x}^2) - u_1(\bar{x}^1 - x_{2,1}, \bar{x}^2 - x_{2,1}) \end{aligned}$$

where the last inequality follows from the assumption that $u_1(y, y)$ is increasing in y (Assumption 2). Hence the assumption that u_1 is decreasing in x_1 implies that $x_{1,1} > x_{2,1}$.

Combining (2) and (4) yields that

$$\begin{aligned} u_2(\bar{x}^1 - x_{1,1}, \bar{x}^2 - x_{2,1} - x_{2,2}) &= u_2(\bar{x}^1, \bar{x}^2 - x_{2,1}) \\ &\geq u_2(\bar{x}^1 - x_{2,1}, \bar{x}^2 - 2x_{2,1}) \\ &\geq u_2(\bar{x}^1 - x_{1,1}, \bar{x}^2 - 2x_{2,1}) \end{aligned} \tag{5}$$

where the second line follows from the assumption that $u_2(y, y)$ is increasing in y (Assumption 2) and the third line follows from $x_{1,1} \geq x_{2,1}$. Expression (5) together with the assumption that u_2 is decreasing in x_2 implies that

$$\begin{aligned}\bar{x}^2 - x_{2,1} - x_{2,2} &\leq \bar{x}^2 - 2x_{2,1} \\ x_{2,2} &\geq x_{2,1}.\end{aligned}$$

The consideration for player two and future rounds is analogous. Thus, investments are increasing over time, i.e., $x_{i,t} \geq x_{i,t-1}$ for either player i and each t where $2 \leq t \leq T-1$. Hence, we can reach the target within at most $\max_i \bar{x}^i / x_{i,1}$ rounds.

ii) For showing that T as implicitly defined in i) is the minimal number of rounds required to implement the target, it suffices to show that relaxing the incentive compatibility constraint for either player tightens subsequent constraints of the other player and also of himself.

First, consider relaxing the incentive compatibility constraint of player two in some round $t-1$, i.e. reducing $x_{2,t-1}$. How does this affect the maximal investment of the player one in round $t-1$, i.e., $x_{1,t-1}$? Denote the sum of future investments as $\hat{x}^i := \bar{x}^i - \sum_{k=1}^{t-1} x_{i,k}$ for $i = 1, 2$. Then $x_{1,t-1}$ is implicitly given by

$$u_1(\hat{x}^1, \hat{x}^2) = u_1(\hat{x}^1 - x_{1,t-1}, \tilde{x}^2 - x_{2,t-1}) - \varepsilon.$$

Interpreting $x_{1,t-1}$ as function of $x_{2,t-1}$, we take implicit derivatives and obtain

$$\frac{\partial x_{1,t-1}}{\partial x_{2,t-1}} = -\frac{\partial u_1 / \partial x^2}{\partial u_1 / \partial x^1} > 0$$

Hence, whenever player two invests less then described by the schedule in i), also player one invests less in that round when all constraints are binding.

Second, consider the effect of relaxing the incentive compatibility constraint of player two in round $t-1$ (holding the behaviour of player one fixed) on this player's maximal investment in round t. Then $x_{2,t}$ is implicitly given by

$$u_2(\hat{x}^1, \hat{x}^2) = u_2(\hat{x}^1 - x_{1,t-1}, \hat{x}^2 - x_{2,t-1} - x_{2,t}) - \varepsilon.$$

Taking implicit derivatives and using Assumption 2 and our assumptions on payoffs we obtain

$$
\begin{aligned}
\frac{\partial x_{2,t}}{\partial x_{2,t-1}} =& -\frac{\partial u_2/\partial x^2 + \partial u_2/\partial x^1 \left(-\frac{\partial u_1/\partial x^2}{\partial u_1/\partial x^1}\right)}{\partial u_2/\partial x^2} \\
=& \frac{\partial u_2/\partial x^1 \cdot \partial u_1/\partial x^2 - \partial u_2/\partial x^2 \cdot \partial u_1/\partial x^1}{\partial u_2/\partial x^2 \cdot \partial u_1/\partial x^1} \\
\geq& \frac{-\partial u_2/\partial x^2 \cdot \partial u_1/\partial x^2 - \partial u_2/\partial x^2 \cdot \partial u_1/\partial x^1}{\partial u_2/\partial x^2 \cdot \partial u_1/\partial x^1} \\
\geq& \frac{\partial u_2/\partial x^2 \cdot \partial u_1/\partial x^1 - \partial u_2/\partial x^2 \cdot \partial u_1/\partial x^1}{\partial u_2/\partial x^2 \cdot \partial u_1/\partial x^1} = 0
\end{aligned}
$$

Hence, whenever player one invests less then described by the schedule in i), this also makes himself invest less in subsequent rounds, increasing the total number of rounds (and analogously for the other player).

Taking these considerations together, the total number of rounds required to implement the target is weakly larger than T whenever we relax a constraint. Therefore, T is the minimal number of rounds required, i.e., an investment schedule with $T' < T$ that implements the target in an ε-SPE does not exist. □

Remark: The investment schedule in Proposition 1 is increasing over time, i.e., $x_{i,t} \geq x_{i,t-1}$ for all $t < T$, $i = 1, 2$. This is shown in the proof of part i) of the proposition.

Indeed, showing that investments are increasing over time is a key building block for proving statement i) of the proposition. The intuition behind this result is as follows. Suppose player two makes an investment $x_{2,t}$ in round t, followed by an investment $x_{1,t}$ of player one. When player two moves again in round $t+1$, he is at least as inclined to make an investment of size $x_{1,t}$ as he was in round t, because now there is more to be gained for him due to Assumption 2.

The key ingredient to proving that the schedule proposed in Proposition 1 is optimal is showing that loosening an incentive compatibility constraint for a certain player in some round tightens both this player's and also the other player's constraint in subsequent rounds, and hence there is less future investment.

Our approach requires that $\varepsilon > 0$. For $\varepsilon = 0$, the ε-SPE degenerates to a SPE, in which the holdup problem replicates itself in the very last round and hence positive investments levels cannot be sustained. What happens as ε becomes arbitrarily close to 0? For the ferryman example, only two rounds are needed regardless how small ε is. If the payoffs of each player are continuous, then as ε tends to zero the number of rounds needed to implement a given target tends to infinity. The formal statement can be found in the appendix (Proposition 7).

Nevertheless, in numerical examples we find that most settings require only a small number of rounds for values of ε that are not excessively small. Following up on the example used above where $u_1(x^1, x^2) = -x^1 + x^2$ and $u_2(x^1, x^2) = 2\sqrt{x^1} - x^2$, 18 rounds suffice to implement the target $(1, 1)$ when $\varepsilon = 0.01$ (in contrast to 50 rounds under constant investments). Seven rounds suffice to implement the target when $\varepsilon = 0.05$ (in contrast to 10 under constant investments; see Section 5 for more numerical examples).

Our takeaway from these examples is twofold. First, the number of rounds required to implement the target might improve substantially when our proposed schedule is used as opposed to constant investments. Moreover, despite the fact that in principle the number of rounds goes to infinity, implementation is typically fast for reasonable values of ε.

3.3 Buyer-seller relationships

The investment schedule described in Proposition 1 implicitly defines the maximal investment that either player is willing to make when he moves. These investments are typically non-linear functions of own and the other player's past (chronologically speaking, future) investments. In this section, we impose additional structure on the payoff functions and show that these additional assumptions result in a very neat characterization of the optimal investments.

More specifically, we restrict payoff functions from the general model as follows. Player two's actions are entirely welfare neutral. This player is referred to as the buyer in this section, and his actions are payments. Social welfare increases only through investments of player one (the seller) with payoffs $u_1(x^1, x^2) = -x^1 + x^2$. Payoffs for the buyer are given by $u_2(x^1, x^2) = g(x^1) - x^2$, where g is referred to as production function. The assumptions made up to now impose that g is differentiable and increasing in x, and that $g(x) - x$ is increasing. In the following, we refer to this as a buyer-seller setting. Notice that the example used above where $u_1(x^1, x^2) = -x^1 + x^2$ and $u_2(x^1, x^2) = 2\sqrt{x^1} - x^2$ falls in this class, where $g(x^1) = 2\sqrt{x^1}$.

Proposition 2. *In buyer-seller settings, the investment schedule* $(x_t)_{t=1}^{T}$ *to implement the target* $\bar{x}$ *in an* ε*-SPE as described in Proposition 1 involves pay-as-you-go for the buyer, i.e., the buyer compensates the seller for his investment in each round* t *where*

$2 \leq t \leq T-1$*:* $x_{1,1} = 2\varepsilon$*,* $x_{2,1} = 1\varepsilon$*,*

$$x_{1,t} = x_{2,t} = g\left(\sum_{i=t-1}^{T} x_{1,i}\right) - g\left(\sum_{i=t}^{T} x_{1,i}\right),$$

for $2 \leq t \leq T-1$*,* $x_{1,T} = \bar{x}^1 - \sum_{i=1}^{T-1} x_{1,i}$ *and* $x_{2,T} = \bar{x}^2 - \sum_{i=1}^{T-1} x_{2,i}$*.*

Proof. See appendix. □

Note that in the investment schedule described in Proposition 2, the seller invests in round t exactly the additional utility for the buyer that is induced by this investment. In that sense, the schedule involves invest-as-you-go as well as pay-as-you-go.

3.4 Discounting in buyer-seller relationships

Even in our main analysis, we account for the fact that payoffs that are further in future are less valuable by describing a schedule that implements targets within a minimal number of rounds. Nevertheless, one may wonder how sensitive our results are to explicitly allowing for discounting.

Unfortunately, the baseline model becomes untractable when incorporating discounting. Instead, we show how our main insights readily generalize for the case with discounting for buyer-seller settings.

We now discount future payoffs with a common discount factor $\delta \in (0,1)$. Present round payoffs accrue immediately. We show that the schedule described in the general model in Proposition 1 remains optimal even with discounting.

Proposition 3. *Consider a buyer-seller game and fix* $\varepsilon > 0$*, a target* $\bar{x}$ *and a discount factor* $\delta \in (0,1)$*.*

i) There exists a T *and an investment schedule* $(x_t)_{t=1}^T$ *such that incentive compatibility constraints (1) are binding for* $t < T$ *for each player and the target* $\bar{x}$ *is implemented in an* ε*-SPE.*

ii) There does not exist a $T' < T$ *such that* $\bar{x}$ *can be implemented in an* ε*-SPE that involves* T' *rounds.*

Proof. Note that since incentive compatibility constraints are binding, no player has an incentive to deviate in any round. Thus in order to proof i), we only need to make sure that eventually the sum of investments equals the target. This is shown as follows.

When the buyer moves in round one it has to hold that

$$g(\bar{x}^1) - \bar{x}^2 = g(\bar{x}^1) - (\bar{x}^2 - x_{2,1}) - \varepsilon$$
$$x_{2,1} = \varepsilon.$$

Similarly, the incentive compatibility constraint in round one for the seller is as follows:

$$(1-\delta)(\bar{x}^2 - x_{2,1}) + \delta \cdot \bar{x}^2 - x_{1,1} = \bar{x}^2 - x_{2,1} - \varepsilon$$
$$x_{1,1} = \delta x_{2,1} + \varepsilon.$$

Note that $x_{1,1}$ is increasing in $x_{2,1}$.

When the buyer moves in round two it has to hold that

$$(1-\delta)\left(g(\bar{x}^1 - x_{1,1}) - (\bar{x}^2 - x_{2,1})\right) + (1-\delta)\delta\left(g(\bar{x}^1) - (\bar{x}^2 - x_{2,1})\right) + \delta^2(g(\bar{x}^1) - \bar{x}^2) =$$
$$g(\bar{x}^1 - x_{1,1}) - (\bar{x}^2 - x_{2,1} - x_{2,2}) - \varepsilon$$
$$x_{2,2} = \delta g(\bar{x}^1) + \varepsilon - \delta g(\bar{x}^1 - x_{1,1}) - \delta^2 x_{2,1}$$

Using $x_{1,1} = \delta x_{2,1} + \varepsilon$ we have that

$$x_{2,2} = \delta g(\bar{x}^1) + \varepsilon - \delta g(\bar{x}^1 - \delta x_{2,1} - \varepsilon) - \delta^2 x_{2,1}$$
$$\frac{\partial x_{2,2}}{\partial x_{2,1}} = \delta^2(g'(\bar{x}^1 - \delta x_{2,1} - \varepsilon) - \bar{x}^2) > 0$$

where the last inequality follows from the assumption that $g(x) - x$ is increasing, i.e., that $g' > 1$. Thus, analogous to the proof of Proposition 1, this establishes ii), i.e., that we cannot gain by loosening incentive compatibility constraints in any round, and hence that the target cannot be implemented with less that T rounds.

Statement i) is proven by the fact that $x_{1,1} \geq \varepsilon$, $x_{2_1} \geq \varepsilon$ and that investments are increasing. Hence we reach the target $\bar{x}$ within at most $\max(\bar{x}^1, \bar{x}^2)/\varepsilon$ rounds. □

Proposition 3 shows that a schedule with binding constraints is still optimal under discounting. Clearly, discounting makes it harder to implement targets, because player's impatience make them more inclined to consume what is available right away instead of waiting for better outcomes in the future. But what about the magnitude of this effect? For instance, for $g(x) = 2\sqrt{x}$ and $\varepsilon = 0.05$, the number of rounds increases from seven to nine as we reduce from δ from 1 to 0.9. The effect becomes stronger for smaller ε though. When $\varepsilon = 0.01$, the number of rounds increases from 18 to 50 (see Section 5 for additional numerical examples).

3.5 Games with at least three players

Maintaining the sequential (alternating moves) structure, we now generalize our main results for settings with $n \geq 3$ players.

The only assumption that we made up to now on the relationship between payoffs and investments is that each player benefits when all players increase their investments equally (Assumption 2). In two-player settings, this suffices to ensure that in round one, player one is as least as incentivized to make an investment as player two was, given that player one anticipates subsequent investment from player two. By Assumption 2, player one prefers a situation in which both invested to a situation in which neither of them invested. This no longer suffices in n-player settings, because Assumption 2 only concerns the effect on payoffs in case *everyone* made an investment, not only a subset of players. Hence the aforementioned argument no longer bites. In particular, player two need no longer be as inclined to make an investment in round one as player three was.

Thus, we are required to make an additional assumption on the *minimal relative utility improvement* Δ defined as

$$\Delta := \min_j \min_k \frac{\partial u_j / \partial x^k}{-\partial u_j / \partial x^j} \; \forall j \neq k. \tag{6}$$

Intuitionally, Δ captures the minimal improvement in utility that any player k can obtain from a one-unit increase in x from each player. As our numerical examples below show, this assumption is typically not very stringent.

In the next proposition, we show that our previous results generalize once Δ is sufficiently high. For further usage, we define $\beta(n)$ as the real root of the following expression:

$$\sum_{j=1}^{n-1} \left(\beta(n) - 1 + \sum_{i:j<i<n} \beta(n) z_i \right) = 0 \tag{7}$$

where z_j is defined recursively as

$$z_j := \beta(n) \left(1 + \sum_{i=j+1}^{n} z_i \right)$$

for $j \in \{1, 2, \ldots, n-1\}$ and $z_n = 0$. Notice that upon expanding the z_i terms, (7) only involves terms of the form $a\beta(n)^b + c(-1)$ for $a > 0$, $b > 0$, $c > 0$. Hence (7) is strictly negative at $\beta(n) = 0$ and strictly increasing in $\beta(n)$, implying that there is a unique real root.

Proposition 4. *Consider a sequential-moves game and fix $\varepsilon > 0$, a target $\bar{x}$ and $n \geq 3$ players. Assume that $\Delta \geq \beta(n)$.*

i) There exists a T and an investment schedule $(x_t)_{t=1}^T$ such that incentive compatibility constraints (1) are binding for $t < T$ for each player and the target $\bar{x}$ is implemented in an ε-SPE.

ii) There does not exist a $T' < T$ such that $\bar{x}$ can be implemented in an ε-SPE that involves T' rounds.

Proof. i) We first show that the target $\bar{x}$ can be implemented in an ε-SPE with constant investment. For each player j, define the scalar x'_j as the maximal investment player j would be willing to make even if everyone else already reached the target, i.e., as the solution to

$$u_j\left(\bar{x}\right) = u_j\left(\left(x^i - x'_i \mathbb{1}_{\{i=j\}}\right)_i\right) - \varepsilon.$$

Notice that for the player moving last this exactly coincides with the condition when there are only two players. Next, define $x' := \min_i x'_i$. By definition of x'_j, this implies that for each player j it holds that

$$u_j\left(\bar{x}\right) - u_j\left(\left(x^i - x' \mathbb{1}_{\{i=j\}}\right)_i\right) + \varepsilon \geq 0, \tag{8}$$

i.e., no player has an incentive to deviate from an investment of x' in the last round, even if no other player invests.

Now suppose that each player invests x' in the last round. This loosens the condition in (8) since

$$u_j\left(\bar{x}\right) - u_j\left(\left(x^i - x' \mathbb{1}_{\{i\geq j\}}\right)_i\right) + \varepsilon \geq u_j\left(\bar{x}\right) - u_j\left(\left(x^i - x' \mathbb{1}_{\{i=j\}}\right)_i\right) + \varepsilon \geq 0$$

by the assumption that payoffs increase in investments of other players and hence neither player deviates in the last round.

Now suppose that each player invests x' in each round t. Again, neither player has an incentive to deviate in round t since

$$u_j\left(\bar{x}\right) - u_j\left(\left(x^i - x' \mathbb{1}_{\{i\geq j\}} - x'(t-1)\right)_i\right) + \varepsilon \geq u_j\left(\bar{x}\right) - u_j\left(\left(x^i - x' \mathbb{1}_{\{i\geq j\}}\right)_i\right) + \varepsilon \geq 0$$
$$u_j\left(\left(x^i - x' \mathbb{1}_{\{i\geq j\}}\right)_i\right) \geq u_j\left(\left(x^i - x' \mathbb{1}_{\{i\geq j\}} - x'(t-1)\right)_i\right)$$

where the last inequality follows from Assumption 2. Hence, using these constant investments, we can reach the objective within at most $\max_j \bar{x}_j / x'$ rounds. As we show in ii), making incentive compatibility constraints binding in each makes players invest even more. Hence we can reach the target with a schedule in which incentive compatibility constraints are binding.

ii) Incentive compatibility of player j in round t requires that

$$u_j\left(\bar{x}\right) \geq u_j\left(\left(\bar{x}^i - \sum_{k=1}^{t-1} x_{i,k} - x_{i,t}1_{\{i\geq j\}}\right)_i\right) - \varepsilon.$$

Now consider changing $x_{n,t}$. Let $x_{j,t}$ as function of $x_{n,t}$ solve

$$u_j\left(\bar{x}\right) = u_j\left(\left(x^i - \sum_{k=1}^{t-1} x_{i,k} - x_{i,t}1_{\{i\geq j\}}\right)_i\right) - \varepsilon.$$

By implicitly differentiating we obtain

$$\frac{\partial u_j}{\partial x^j}\left(-\frac{\partial x_{j,t}}{\partial x_{n,t}}\right) + \sum_{i:j<i<n} \frac{\partial u_j}{\partial x^i}\left(-\frac{\partial x_{i,t}}{\partial x_{n,t}}\right) + \frac{\partial u_j}{\partial x^n}\left(-1\right) = 0$$

so

$$\frac{\partial x_{j,t}}{\partial x_{n,t}} = \frac{\frac{\partial u_j}{\partial x^1} + \sum_{i:j<i<n} \frac{\partial u_j}{\partial x^i}\left(\frac{\partial x_{i,t}}{\partial x_{n,t}}\right)}{\left(-\frac{\partial u_j}{\partial x^j}\right)}.$$

Note that our assumptions imply that $\frac{\partial x_{j,t}}{\partial x_{n,t}} \geq 0$.

Now suppose that

$$\frac{\frac{\partial u_j}{\partial x^k}}{\left(-\frac{\partial u_j}{\partial x^j}\right)} \geq \beta$$

for all $j \neq k$. Then

$$\frac{\partial x_{j,t}}{\partial x_{n,t}} = \frac{\frac{\partial u_j}{\partial x^n} + \sum_{i:j<i<n} \frac{\partial u_j}{\partial x^i}\left(\frac{\partial x_{i,t}}{\partial x_{n,t}}\right)}{\left(-\frac{\partial u_j}{\partial x^j}\right)} \geq \beta + \beta \sum_{i:j<i<n} \frac{\partial x_{i,t}}{\partial x_{n,t}}$$

and hence $\frac{\partial x_{n-1,t}}{\partial x_{n,t}} \geq \beta$, $\frac{\partial x_{n-2,t}}{\partial x_{n,t}} \geq \beta + \beta^2$, and more generally $\frac{\partial x_{j,t}}{\partial x_{n,t}} \geq z_j$ where $z_j = \beta\left(1 + \sum_{i=j+1}^{n} z_i\right)$.

Now consider

$$\frac{\partial u_n}{\partial x^n}\left(-1 - \frac{\partial x_{n,t+1}}{\partial x_{n,t}}\right) + \sum_{j<n} \frac{\partial u_n}{\partial x^j}\left(-\frac{\partial x_{j,t}}{\partial x_{n,t}}\right) = 0$$

so

$$\frac{\partial x_{n,t+1}}{\partial x_{n,t}} = \frac{\frac{\partial u_n}{\partial x^n} + \sum_{j<n} \frac{\partial u_n}{\partial x^j} \frac{\partial x_{j,t}}{\partial x_{n,t}}}{\left(-\frac{\partial u_n}{\partial x^n}\right)}$$

By our assumption that $\sum_{j=1}^{n} \frac{\partial u_k}{\partial x^j} \geq 0$ for each k (equal increase in contribution from each other player helps each player) follows that

$$\begin{aligned}
\frac{\partial x_{n,t+1}}{\partial x_{n,t}} &\geq \frac{\sum_{j<n} \frac{\partial u_1}{\partial x^j}\left(\frac{\partial x_{j,t}}{\partial x_{n,t}} - 1\right)}{\left(-\frac{\partial u_n}{\partial x^n}\right)} \\
&= \frac{\sum_{j<n} \frac{\partial u_n}{\partial x^j}\left(\frac{\frac{\partial u_j}{\partial x^n} + \sum_{i:j<i<n}\left(\frac{\partial u_j}{\partial x^i}\left(\frac{\partial x_{i,t}}{\partial x_{n,t}}\right)\right)}{\left(-\frac{\partial u_j}{\partial x^j}\right)} - 1\right)}{-\frac{\partial u_n}{\partial x^n}} \\
&= \sum_{j<n} \frac{\partial u_n}{\partial x^j}\left(\frac{\frac{\partial u_j}{\partial x^1} + \frac{\partial u_j}{\partial x^j} + \sum_{i:j<i<n} \frac{\partial u_j}{\partial x^i}\left(\frac{\partial x_{i,t}}{\partial x_{n,t}}\right)}{\frac{\partial u_j}{\partial x^j}\frac{\partial u_n}{\partial x^n}}\right).
\end{aligned}$$

Now we use the assumption that $\frac{\partial x_{j,t}}{\partial x_{n,t}} \geq z_j$ to obtain

$$\frac{\partial x_{n,t+1}}{\partial x_{n,t}} \geq \beta \sum_{j=1}^{n-1}\left(\beta - 1 + \sum_{i:j<i<n} \beta z_i\right). \tag{9}$$

The sum in the right-hand side of (9) is exactly the definition of $\beta(n)$. As we argued above, this expression is increasing in β and has a unique real root. Hence whenever $\Delta \geq \beta(n)$, (9) is positive. Thus own future investments are increasing in own investments, guaranteeing that $\bar{x}$ cannot be implemented with less than T rounds. □

Proposition 4 gives a lower bound on the minimal relative utility improvement Δ such that our approach is guaranteed to be optimal in the sense of requiring the minimal number of rounds. Other than the additional requirement that Δ is sufficiently high, the proposition is a generalization of Proposition 1. The intuition behind the proof is analogous.

As an example for why the additional condition is needed, consider the following three-player symmetric linear public goods game. The game is played sequentially over multiple rounds, with the following payoff function for each player i:

$$u_i(x_i, x_{-i}) = \alpha\left(\sum_{k=1}^{3} x_k\right) - x_i$$

where the marginal per capita return α satisfies $1/3 < \alpha < 1$. This gives the game our standard properties, in particular $\partial u_i/\partial x_i = \alpha - 1$, $\partial u_i/\partial x_j = \alpha$, and both social welfare and individual payoffs increase when all players do the same. We normalize such that the social optimum is attained when all players invest 1, which is also our target.

In this example, for (9) to hold requires that

$$\beta(\beta-1)+\beta(\beta-1+z_1\beta)>0$$

where $z_1=\beta(1+\beta)$. Hence we need that

$$\beta(\beta^3+\beta^2+2\beta-2)\geq 0$$
$$\beta \gtrapprox 0.65.$$

For a linear public goods game we have that $\frac{\alpha}{1-\alpha}=\beta$ and hence this requires that $\alpha \gtrapprox 0.394$. Note that this is obtained generically, i.e., the β is calculated independently of the actual payoff functions and of ε and only a function of n.

How much are we loosing by the generic assumption on Δ? Alternatively speaking, how close is this bound from the exact bound? To answer these questions, we explicitly solve the three-player public goods game

Denote by $x_{i,t}$ the investment of player i in round t.

So for the player moving last we have that $x_{3,1}$ is given by:

$$u_3(1,1,1)=u_3(1,1,1-x_{3,1})-\varepsilon$$
$$x_{3,1}=\frac{\varepsilon}{1-\alpha}$$

Then for player two we have that

$$u_2(1,1,1)=u_2(1,1-x_{2,1},1-x_{3,1})-\varepsilon$$
$$x_{2,1}=\frac{x_{3,1}\alpha+\varepsilon}{1-\alpha}$$

so $x_{2,1}$ is increasing in $x_{3,1}$, as desired.

And for player one we have that

$$u_1(1,1,1)=u_1(1-x_{1,1},1-x_{2,1},1-x_{3,1})-\varepsilon$$
$$x_{1,1}=\frac{\alpha(x_{2,1}+x_{3,1})+\varepsilon}{1-\alpha}$$

so $x_{1,1}$ is also increasing in $x_{3,1}$, as desired.

When player three moves again in round two, his maximal investment is given by

$$u_3(1,1,1)=u_3(1-x_{1,1},1-x_{1,2},1-x_{3,1}-x_{3,2})-\varepsilon$$
$$x_{3,2}=\frac{-x_{3,1}(1-\alpha)+\alpha(x_{1,1}+x_{2,1})+\varepsilon}{1-\alpha}$$

Upon expanding that expression we find that

$$x_{3,2} = \frac{x_{3,1}(1 - 3\alpha + \alpha^2) - \varepsilon}{(\alpha - 1)^3}$$

which decreases in $x_{3,1}$ for α sufficiently close to $1/3$ (the exact cut-off point is $(3 - \sqrt{5})/2 \approx 0.382$). Since $x_{3,2}$ may be decreasing in $x_{2,1}$, our approach with making all constraints in all rounds binding need not be optimal any more when using only Assumption 2. The exact cut-off α of 0.382 is very close to our 0.394 generically computed above.

4 Simultaneous investments with symmetric players

In this section, we consider simultaneous move games with $n \geq 2$ players while keeping the remainder of the model as in the previous sections. For simplicity we restrict attention to symmetric games, i.e., $u(x) := u_j(x) = u_i(x)$ and $x_{i,t} = x_{j,t}$ for all $i \neq j$. As we will show, the symmetry makes our approach even simpler, because all players now always face the same constraints. For simultaneous moves games, using only Assumption 2 suffices in order to show optimality of our investment schedule. In particular, we are no longer required to make an additional assumption on Δ as was the case for n-player sequential games. The reason is the following. In simultaneous games, in round $t+1$ all players are at least as inclined to make an investment as they were in round t. This follows from the observation that everyone made an equal investment in round t, which is beneficial for either player by Assumption 2.

Proposition 5. *Consider a simultaneous and symmetric game and fix $\varepsilon > 0$, a target $\bar{x}$ and $n \geq 2$ players.*

i) There exists a T and an investment schedule $(x_t)_{t=1}^{T}$ such that incentive compatibility constraints (1) are binding for $t < T$ for each player and the target $\bar{x}$ is implemented in an ε-SPE.

ii) There does not exist a $T' < T$ such that $\bar{x}$ can be implemented in an ε-SPE that involves T' rounds.

Proof. The maximal investment x_1 in round one satisfies

$$u(\bar{x}^1, \bar{x}^2, \ldots, \bar{x}^n) = u(\bar{x}^1 - x_1, \bar{x}^2, \ldots, \bar{x}^n) - \varepsilon$$

and hence the maximal investment x_2 in round two satisfies

$$u(\bar{x}^1, \bar{x}^2, \dots, \bar{x}^n) = u(\bar{x}^1 - x_1 - x_2, \bar{x}^2, \dots, \bar{x}^n) - \varepsilon.$$

It holds that $x_2 > x_1$ since

$$u(\bar{x}^1 - x_1 - x_2, \bar{x}^2, \dots, \bar{x}^n) < u(\bar{x}^1 - x_1, \bar{x}^2, \dots, \bar{x}^n)$$

by Assumption 2. Thus, investments are increasing over time. This proves that we reach the target $\bar{x}$ within at most $\bar{x}/x_1$ rounds. Since incentive compatibility constraints are binding by construction, this is an ε-SPE.

Statement ii) is proven by the observation that loosening constraints in one round does not yield higher investments in other rounds:

$$\frac{\partial x_2}{\partial x_1} = -\frac{\partial u_2/\partial x^1 + \partial u_2/\partial x^2 + \partial u_2/\partial x^3 + \dots}{\partial u_2/\partial x^1} \geq 0$$

where the last inequality follows from the fact that the numerator is positive by Assumption 2. □

The intuition behind proving Proposition 5 is analogous to Proposition 1.

5 Numerical examples

We now present various numerical examples, some of which were also already mentioned in the main text. In Table 1 we consider two-player buyer-seller games, without discounting (see Section 3.3) and with discounting (see Section 3.4). We show the number of rounds required to implement the socially efficient investment levels of $(1, 1)$, and also the number of rounds required under constant investments. For instance, for the production function $g(x) = 2\sqrt{x}$ and $\varepsilon = 0.05$, our approach requires only seven rounds, whereas under constant investments, 10 are required. The difference is even more severe when ε gets smaller. Our approach is also very robust under moderate levels of discounting.

In Figure 1, we present the investments in each round of an ε-SPE of a buyer-seller game for $g(x) = 2\sqrt{x}$ and different values of ε. The smaller ε, the longer it takes until substantial investments are undertaken, contributing to a longer total duration.

We also present the minimal number of rounds required to make every player invest 1 in linear public goods games, as a function of ε and marginal per capita return α. Table

2 considers a sequential two-player game, and Table 3 a sequential three-player game. In Table 4 we have a simultaneous three-player game. Figure 2 shows the investments in the simultaneous three-player linear public goods game for $\alpha = 1/2$ and different levels of ε. The number of rounds required in the simultaneous game is higher, but the total number of moves is actually smaller because three players move simultaneously.

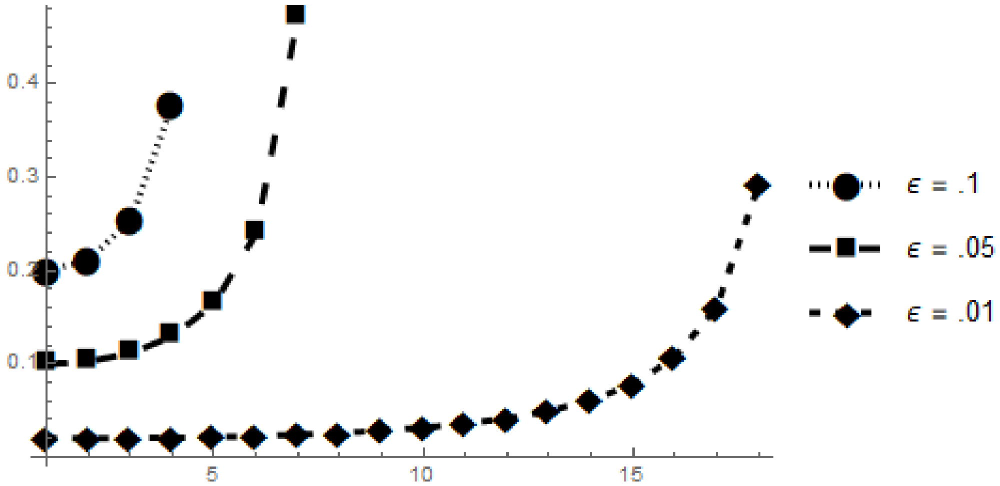

Figure 1: This figure shows investments in an ε-SPE for $g(x) = 2\sqrt{x}$, given ε, over time. Note that round one is the last round.

		$\min T$, $\delta = 1$			$\delta = .99$		$\delta = .95$		$\delta = .9$	
	ε	$2\sqrt{x}$	$1.1x$	$1.1 \cdot \mathbb{1}_{\{x \geq 1\}}$	$2\sqrt{x}$	$1.1x$	$2\sqrt{x}$	$1.1x$	$2\sqrt{x}$	$1.1x$
ε-SPE	.01	18	19	2	22	26	43	39	50	45
	.05	7	8	2	7	8	8	9	9	10
	.1	4	5	2	4	5	5	5	5	5
Const. ε-SPE	.01	50	46	2						
	.05	10	10	2						
	.1	5	5	2						

Table 1: The upper panel of this table shows the minimal number of rounds T required to implement the target $(1, 1)$ in buyer-seller settings, given different production functions g and discounting factors δ. The lower panel show the number of rounds required using constant investments.

ε.

	α		
ε	.51	.75	.9
.01	15	3	2
.05	5	2	2
.1	3	2	1

Table 2: This table shows the number of rounds required to implement the target $(1, 1)$ in a two-player sequential public goods game, given marginal per capita return α and ε.

	α		
ε	.34	.5	.75
.01	19	4	2
.05	7	3	2
.1	4	2	2

Table 3: This table shows the number of rounds required to implement the target $(1, 1, 1)$ in a three-player sequential public goods game, given marginal per capita return α and ε.

	α		
ε	.34	.5	.75
.01	37	6	3
.05	12	4	2
.1	7	3	2

Table 4: This table shows the number of rounds required to implement the target $(1, 1, 1)$ in a three-player simultaneous public goods game, given marginal per capita return α and ε.

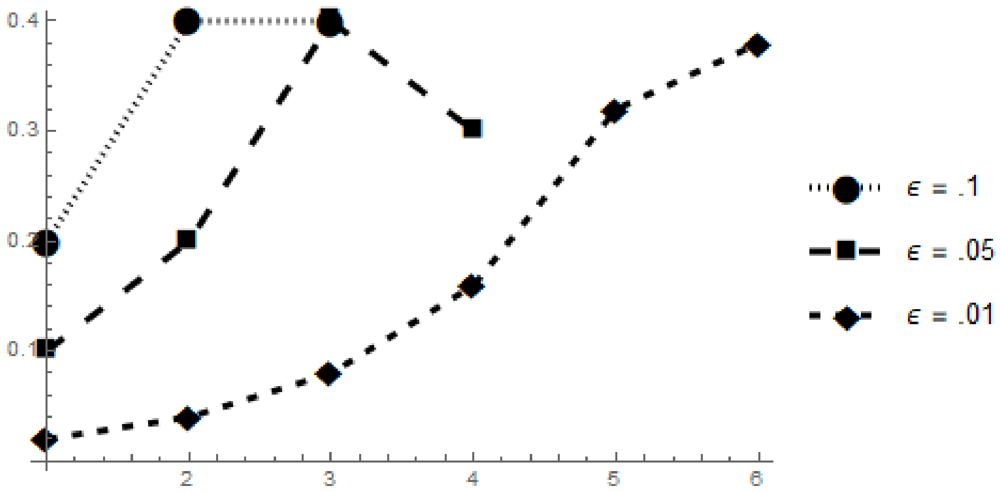

Figure 2: This figure shows the investments in a three-player simultaneous public goods game, given marginal per capita return $\alpha = 1/2$ and ε.

6 Extensions

6.1 Social preferences

In this section we briefly discuss how social preferences can be used to help implement mutually beneficial investments. We continue to split up the investments as one ingredient for reducing incentives to deviate. However, we return to the classic concept of subgame perfection, thus assume that players even care about peanuts. The problem of incentivizing cooperative investments originates in the last round. Investments are costly and hence will not be chosen if a player acts in her own best interest. However, if a player also cares to some degree about the payoffs of the other player then incentives to cooperate can change. In the following we consider a form of conditional social preference relations where a player also cares about the well being of the other if the other has made the desired investments. Very similar arguments apply in the case that each player gains an additional value when the target investment levels are reached.

We start by illustrating the impact of these conditional social preferences in the simplest setting in which each player moves only once. Recall that investments are costly to the player making the investment but beneficial to the other. Player two as last mover will choose the desired investment even if this involves costs to her if she cares sufficiently about the wellbeing of the other player. In particular, if her investment involves negligible costs then she needs to only care a bit about the other player in order to incentivize her to invest. What about player one? If player one deviates then player two turns self interested and chooses inaction. So player one will not deviate either. So both players can be incentivized to invest as desired as part of a subgame perfect equilibrium provided player two cares sufficiently about the well being of player one.

Formally, we are considering utility functions U_i such that $U_i(a) = u_i(a) + \lambda_i \mathbb{1}_{\{a=\bar{a}\}} u_{-i}(a)$ where $\lambda_i > 0$ is given[3]. The sufficient condition for implementing the desired investments $\bar{a}$ as part of a subgame perfect equilibrium is that $u_2(\bar{a}) + \lambda_2 u_1(\bar{a}) \geq u_2(\bar{a}_1, n_2) + \lambda_2 u_1(\bar{a}, n_2)$ which means that $u_2(\bar{a}_1, n_2) - u_2(\bar{a}) \leq \lambda_2 (u_1(\bar{a}) - u_1(\bar{a}, n_2))$. So the care of player two for the benefits for player one has to outweigh the costs to player two.

When the costs of player two are too large then additional rounds of investments need to be added, analogously to our previous analysis. Conditions are very similar to our previous setting. For instance, consider the case with three rounds of investment where player two prior moves prior to player one who is then followed by player two. Replace the two equations involving ε by $U_2(\bar{a}) \geq U_2(\bar{a}_1, (\bar{a}_{21}, n_2))$ and $U_1(\bar{a}) \geq U_1(n_1, (\bar{a}_{21}, n_2))$. This works for the ferryman example if the cost of letting the passenger off the boat is smaller than how much the ferryman cares that the passenger made it to the other side.

Note that only a small degree of social preferences is needed if the payoff of player one has a discontinuous upwards jump as player two increases her investment towards her target level. In fact, it is the social preferences of the last mover that plays the role.

More general solutions are easily derived using the methodology of the previous sections.

6.2 Discounting and costly splitting

To simplify exposition we have assumed throughout that players are infinitely patient and there is no cost of splitting up investments. We now briefly comment why our results continue to hold if players are very patient and costs of splitting up investments are very small.

Consider first impatience and assume that players discount future payoffs with discount

[3] In the alternative scenario where both players have an added value to reaching the desired investments we would have $U_i(a) = u_i(a) + \lambda_i \mathbb{1}_{\{a=\bar{a}\}}$ for some $\lambda_i > 0$, i.e., the utility derived from playing the equilibrium action is independent of the actual utility of the other player.

factor $\delta < 1$. We now argue that our results continue to hold if δ is sufficiently large. Consider first the ability to reach the target investment levels and consider the setting in which we showed that this is possible when $\delta = 1$. Now decrease ε slightly and follow our construction for that the lower value of ε. Then if ε is increased back to its original level then all incentive compatability constraints hold with strict inequality. They will continue to hold if we move to the case of discounting, as long as the discount factor δ is sufficiently large. So the target levels of investment can be reached under sufficiently patient players.

One expects that impatient players require equally or more rounds to implement a given target. Note however that it follows from the arguments made above that at most one more round of investment is needed when δ is sufficiently large.

We now comment on minimal number of rounds needed to implement a given target. We do this for the case where the minimal number of rounds to implement the target under $\delta = 1$ does not change if ε is slightly reduced. Note that the situation in which any slight reduction of ε leads to a strict increase in T under $\delta = 1$ is a knife edge case. Outside of this knife edge case, our construction above shows that sufficiently high discounting cannot increase the number of rounds needed. We now show why it also cannot strictly decrease the number of rounds. This is because any solution for $\delta < 1$ will also be feasible for $\delta = 1$ if ε is slightly increased. Small changes in ε will not change our solution under $\delta = 1$, as we are not looking at the knife edge case. Hence the solutions for $\delta < 1$ cannot yield less rounds when δ is very large.

Note that the above arguments do not replicate one nice feature from the $\delta = 1$ setting. They do not show that one can find the shortest sequence of investments by letting the incentive compatability constraints bind but require a different line of reasoning.

Consider now the situation where splitting up the target investment into smaller investments is costly (or equivalently, adds inefficiencies). Note that introducing very small costs is very similar to reducing δ slightly when $\delta = 1$. In fact, all the previous arguments apply to show that our results on implementation and minimality extend to the case where costs of splitting are sufficiently small.

6.3 Prisoners' dilemma

In this section we illustrate how our insights also allow to support cooperative outcomes in finitely repeated games. This is done by adjusting the time between interactions. Shorter time means smaller payoffs and hence smaller benefits from deviating. To keep the exposition simple we only present the approach for a specific Prisoners' Dilemma.

Consider a Prisoners' Dilemma with the following payoffs:

	C	D
C	$3,3$	$0,5$
D	$5,0$	$1,1$

and assume that total time 1 is divided into T rounds, where round t lasts x_t, so $x_t \geq 0$ and $\sum x_t = 1$. Then the payoff of choosing a_t in round t is $\sum_{t=1}^{T} x_t u_i(a_t)$.

We wish to implement cooperation using grim trigger strategies, so choice of C in the first round, and later if and only if both players chose the same action in past.

In the last round we need that

$$x_T u(C,C) \geq x_T u(D,C) - \varepsilon$$
$$\varepsilon = 2x_T$$
$$x_T = \frac{1}{2}\varepsilon.$$

Similarly, in every round k we need that

$$\sum_{t=k}^{T} x_t u(C,C) \geq x_k u(D,C) + \sum_{t=k+1}^{T} x_t u(D,D) - \varepsilon$$

and hence x_k needs to be such that

$$3\sum_{t=k}^{T} x_t = 5x_k + \sum_{t=k+1}^{T} x_t - \varepsilon.$$

Let $z_k = \sum_{t=k}^{T} x_t$, then

$$3z_k = 5(z_k - z_{k+1}) + z_{k+1} - \varepsilon,$$

so $z_k = 2z_{k+1} + \frac{1}{2}\varepsilon$. Hence $z_{T-1} = \frac{3}{2}\varepsilon$, $z_{T-2} = \frac{7}{2}\varepsilon$, $z_{T-3} = \frac{15}{2}\varepsilon$ and generically $z_k = \frac{2^{T-k+1}-1}{2}\varepsilon$ as $\frac{2^{T-k+1}-1}{2} = 2\frac{2^{T-k}-1}{2} + \frac{1}{2}$. Therefore

$$x_k = z_k - z_{k+1} = \frac{2^{T-k+1}-1}{2}\varepsilon - \frac{2^{T-k}-1}{2}\varepsilon = 2^{T-k-1}\varepsilon$$

and we require that

$$1 = \sum x_k = \sum_{k=1}^{T} 2^{T-k-1}\varepsilon = \frac{1}{2}\left(2^T - 1\right)\varepsilon$$

and therefore

$$T = \left\lceil \frac{\ln \frac{\varepsilon+2}{\varepsilon}}{\ln 2} \right\rceil .$$

Thus, for instance for $\varepsilon = \frac{1}{100}$ we have that $T = \lceil 7.65 \rceil = 8$, and for $\varepsilon = \frac{1}{20}$, $T = \lceil 5.36 \rceil = 6$.

7 Conclusion

In this paper we introduce a general and novel method for implementing cooperative solutions within finite horizons. It requires the ability to split up decisions and perfect information about past choices. Its justifications rests on the choice of a solution concept that presumes there will be no deviations if gains are negligible. The underlying payoffs have to satisfy several conditions. Inaction has to be preferred to any other choice. Choices have to benefit others, and via their reaction also oneself. We focus on the case where players are infinitely patient and where there are no costs of splitting up decisions. As shown in the previous section, all results extend to the setting where players are sufficiently patient and costs of splitting up investments are sufficiently small.

We position this paper as a mix between a normative and a positive approach. From the normative viewpoint we implicitly object to a methodology that rules out outcomes just because very small gains can be made. We provide many reasons why such gains will be not undertaken, thereby justifying the choice of ε-SPE as solution concept. We even show how a particular type of social preferences will generate very similar results. The other aspect of our method, the splitting up of the decision, is used to reduce incentives to deviate. This identifies the similarity to Pitchford and Snyder (2004) albeit we use this design feature in a different context. From the descriptive viewpoint we clarify among other things a role of reference letters and payment on delivery practices. We also direct attention to decisions and interactions that appear late in a relationship that seem unimportant and yet whose existence helps incentivizing the earlier more substantial investments. The simplicity of the analysis makes these insights approachable to wider audiences, insights that are easily used in teaching strategic decision making and game theory.

It is difficult to study the implementation of cooperative outcomes under finite horizons without also connecting to finitely repeated games. We touch on this topic by including the Prisoners' Dilemma with changing frequency of interaction and show how we can implement cooperation. Future research should consider the general setting of a finitely repeated game. One can imagine various models for how the time before interacting again is determined. For instance, this time can be part of a strategy where the minimal time specified is the duration. Alternatively one might imagine that each player commits to how long she will choose the strategy. Both of these scenarios can be used to implement mutual cooperation.

Our paper also initiates a more applied research area. How to optimally split up decisions when an entire project involves many different kinds of tasks. More drastically, how to create or add decisions to provide incentives. Note that the previous research has not touched upon either of these topics, a literature that deals with a single dimensional investment. We touch upon this aspect when we verbally present the mechanisms and challenges in Section 2. For instance, one can imagine that the landlord introduces the magnet as a way to keep him incentivized. Alternatively, for a landlord with good, but unobservable, intentions it can be a signal to the tenant that he will try hard, regardless of his intrinsic intentions. For a different example, consider the shipment of a bicycle. Given the mechanisms identified in this paper, the seller should look for a part of the bicycle that is cheap to ship but essential and difficult to replace when wanting to use the bicycle. This can be the joint connecting the handle bar to the rest of the body. The idea is then for the seller to ship the bicycle without this part, for the buyer to pay the entire price and then for the seller to send the missing part.

One general principle that our paper points to is that the game is never over. Just because substantial interactions cease does not mean that there are not additional smaller business exchanges between the parties, additional small business that is used to incentivize earlier larger business. These small businesses may be incentivized by even smaller ones, until reaching the dimensions of peanuts where cooperation for the sake of cooperating can be valid despite (small) gains from defecting.

Appendix

A Proofs

Proposition 6. *Consider a sequential two-player game and fix $\varepsilon > 0$ and a target $\bar{x}$. There exists $c > 0$ and T such that using constant investments $x_{i,t} = c$ for each player i in each round $t < T$, the target can be implemented in an ε-SPE.*

Proof. Let $x_{2,1}$ be the solution to

$$u_2(\bar{x}^1, \bar{x}^2) = u_2(\bar{x}^1, \bar{x}^2 - x_{2,1}) - \varepsilon$$

which guarantees that player two does not deviate in round one. Let $c = x_{2,1}$.

Suppose we also set $x_{1,1} = c$, then player one does not deviate since we have that

$$u_1(\bar{x}^1, \bar{x}^2) > u_1(\bar{x}^1 - c, \bar{x}^2 - c) > u_1(\bar{x}^1 - c, \bar{x}^2 - c) - \varepsilon$$

where the first inequality follows from the assumption that $u_i(y, y)$ is increasing in y (Assumption 2).

In round two, we need to ensure that player two does not deviate from an additional investment $x_{2,2} = c$. This holds since

$$u_2(\bar{x}^1 - c, \bar{x}^2 - 2c) - \varepsilon < u_2(\bar{x}^1 c, \bar{x}^2 - c) - \varepsilon = u_2(\bar{x}^1, \bar{x}^2).$$

Similarly, no player has an incentive to deviate in any other round. Hence we reach $(\bar{x}^1, \bar{x}^2)$ within at most $T = \max(\bar{x}^1, \bar{x}^2)/c$ rounds. □

Proposition 7. *Consider a sequential two-player game and fix $\varepsilon > 0$ and a target $\bar{x}$. As ε goes to 0, the number of rounds T required to implement the target $(\bar{x}^1, \bar{x}^2)$ in an ε-SPE goes to infinity.*

Proof. As we showed in Proposition 1, the investment schedule implementing a target $(\bar{x}^1, \bar{x}^2)$ within minimal rounds requires that incentive compatibility constraints are binding in all rounds. Hence, in round one we have that investments $x_{2,1}$ are given by the solution to

$$u_2(\bar{x}^1, \bar{x}^2) = u_2(\bar{x}^1, \bar{x}^2 - x_{2,1}) - \varepsilon$$

which goes to 0 as ε goes to 0.

Given that, we also have that $x_{1,1}$ as the solution to

$$u_1(\bar{x}^1, \bar{x}^2) = u_1(\bar{x}^1 - x_{1,1}, \bar{x}^2 - x_{2,1}) - \varepsilon$$

goes to 0 as ε goes to 0. Similarly, all other investments go to 0 as ε goes to 0. Hence the number of rounds required to implement the desired outcome $(\bar{x}^1, \bar{x}^2)$ goes to infinity as ε goes to 0. □

Proof of Proposition 2. Incentive compatibility (1) for the buyer in the last round requires that

$$\begin{aligned} u_2(\bar{x}^1, \bar{x}^2) &= u_2(\bar{x}^1, \bar{x}^2 - x_{2,1}) - \varepsilon \\ x_{2,1} &= \varepsilon \end{aligned}$$

and for the seller that

$$\begin{aligned} u_1(\bar{x}^1, \bar{x}^2) &= u_2(\bar{x}^1 - x_{2,1}, \bar{x}^2 - x_{1,1}) - \varepsilon \\ x_{1,1} &= 2\varepsilon. \end{aligned}$$

The constraint for the buyer in each round t where $2 \leq t \leq T-1$ simplifies to

$$x_{2,t} = g(\bar{x}^1) - g\left(\bar{x}^1 - \sum_{k=1}^{t} x_{1,k}\right) - \sum_{k=1}^{t-1} x_{2,k} + \varepsilon$$

and for the seller to

$$\begin{aligned} x_{1,t} &= \sum_{k=1}^{t} x_{2,k} - \sum_{k=1}^{t-1} x_{1,k} + \varepsilon \\ &= \varepsilon + \sum_{k=2}^{t-1} x_{2,k} + x_{1,t} - \sum_{k=2}^{t-1} x_{1,k} - 2\varepsilon + \varepsilon \\ &= x_{2,t}. \end{aligned}$$

Thus, payments equal investments for every t where $2 \leq t \leq T-1$. The round T investments are then set such that the total sum of investments equals $\bar{x}$. Writing the investments recursively we obtain

$$x_{1,t} = g\left(\bar{x}^1 - \sum_{i=1}^{t-2} x_{1,i}\right) - g\left(\bar{x}^1 - \sum_{i=1}^{t-1} x_{1,i}\right)$$

and since $\bar{x}^1 = \sum_{i=1}^{T} x_1, i$, we can write this as

$$x_{1,t} = g\left(\sum_{i=t-1}^{T} x_{1,i}\right) - g\left(\sum_{i=t}^{T} x_{1,i}\right).$$

□

References

Admati, A. R. and Perry, M. (1991). Joint projects without commitment. *The Review of Economic Studies*, 58(2):259–276.

Barlo, M. and Dalkiran, N. A. (2009). Epsilon-nash implementation. *Economics Letters*, 102(1):36–38.

Baye, M. R. and Morgan, J. (2004). Price dispersion in the lab and on the internet: Theory and evidence. *RAND Journal of Economics*, pages 449–466.

Che, Y.-K. and Sákovics, J. (2004). A dynamic theory of holdup. *Econometrica*, 72(4):1063–1103.

Compte, O. and Jehiel, P. (2004). Gradualism in bargaining and contribution games. *The Review of Economic Studies*, 71(4):975–1000.

de Burgh, C. (1982). Don't pay the ferryman. The Getaway, A&M.

Dixit, A. K. and Nalebuff, B. J. (1993). *Thinking strategically: The competitive edge in business, politics, and everyday life*. WW Norton & Company.

Dulleck, U., Kerschbamer, R., and Sutter, M. (2011). The economics of credence goods: An experiment on the role of liability, verifiability, reputation, and competition. *The American Economic Review*, 101(2):526–555.

Lockwood, B. and Thomas, J. P. (2002). Gradualism and irreversibility. *The Review of Economic Studies*, 69(2):339–356.

Mailath, G. J., Postlewaite, A., and Samuelson, L. (2005). Contemporaneous perfect epsilon-equilibria. *Games and Economic Behavior*, 53(1):126–140.

Milgrom, P. (2010). Simplified mechanisms with an application to sponsored-search auctions. *Games and Economic Behavior*, 70(1):62–70.

Nardo, D. (2002). *The Greenhaven Encyclopedias Of - Greek and Roman Mythology*. Greenhaven Press.

Palfrey, T. R. and Prisbrey, J. E. (1996). Altuism, reputation and noise in linear public goods experiments. *Journal of Public Economics*, 61(3):409–427.

Pitchford, R. and Snyder, C. M. (2004). A solution to the hold-up problem involving gradual investment. *Journal of Economic Theory*, 114(1):88–103.

Radner, R. et al. (1980). Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives. *Journal of Economic Theory*, 22(2):136–154.

Shapiro, C. (1982). Consumer information, product quality, and seller reputation. *The Bell Journal of Economics*, pages 20–35.