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

Simon Martin* Karl H. Schlag**

April 6, 2020

Abstract

When should one pay the ferryman? When to pay for the delivery of a good, and how should one invest in a public good if there is a single transaction and institutions are costly? We show how to solve the hold-up problem. The idea is to appropriately split up the desired total contribution into several separate contributions that are made in sequence, with each party threatening to discontinue if others deviate. Our solution concept is based on subgame perfection, where players are either selfish and do not care about very small gains or are socially minded as long as there is no deviation.

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

JEL Codes: D23, C72, L14

[◊]We would like to thank Bernhard Kasberger, Matthijs Wildenbeest, participants of the VGSE Micro Seminar and several workshops and conferences, as well as the Editor Pietro Ortoleva, the Associate Editor and the referees of this submission for their useful comments and suggestions.

*Düsseldorf Institute for Competition Economics (DICE), University of Düsseldorf, simon.martin@dice.hhu.de

**University of Vienna, 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 named 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 has no incentive to continue the passage once he received your payment. 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 is no incentive 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 players be convinced to contribute to a public good when it is a dominant strategy to abstain from such a contribution?

The common feature of these examples is that investments are irreversible and own investments are costly. If there are no possibilities for commitment, desirable allocations cannot be implemented in equilibrium. Thus, there is a hold-up problem. In our examples above, the prediction is that the passenger will not be transported, trade will not take place and the public good will not be provided.

There are different ways to ensure that individuals make the desirable contributions. Institutions can help by making contracts enforceable and allowing allocation and transfer of property rights to better align incentives. 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

[1]For a theoretical contribution, see Shapiro (1982). For experimental investigations, see Dulleck et al. (2011) and Palfrey and Prisbrey (1996).

a buyer and seller is uncertain and potentially never ending, one can overcome incentive problems by splitting the entire investment into smaller parts. However, this solution is not practical as one has to split the transaction into infinitely many parts and the entire good is never completely delivered. The other possibilities mentioned above also have substantial drawbacks. They limit the ability to trade in modern market economies defined by globalization, decentralization and the need for flexibility. Institutions are costly and often not available. Property rights typically only partially mitigate the problem (Grossman and Hart, 1986). Potential trading partners are constantly changing.

In this paper, we propose a novel solution to the hold-up problem in finite time. Without institutions and without the need to condition on other interactions, 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 provides insights into existing practices and qualifies how to better design interactions. Most importantly, our method is particularly simple to implement. We implement the desired allocation by splitting up the total contribution (e.g., investment or payment) of each player into finitely many smaller contributions that are made in sequence. Non-compliance is punished by the termination of the relationship. For our implementation we consider preferences that are slightly different to the standard selfish prototype. A player deviates only if this deviation leads to a substantially higher payoff. This is formally captured by the concept of ε-SPE (Mailath et al., 2005).[2] The parameter ε captures how much is needed to trigger a deviation. In a separate section, we argue why these preferences seem more natural in many cases than the standard ones. Importantly, it suffices that a single player has these preferences in order to make our approach applicable. In a separate section we show that our findings similarly hold under social preferences in which a player cares to some degree about the payoff of others as long as no one has deviated.

With the novel combination of these two ingredients, splitting up the decision and solving with ε-SPE, it is intuitive that the hold-up problem can be solved. Split the total contribution into many very small, possibly equally-sized, contributions and threaten to stop contributing if some player previously did not make the assigned contribution. The last contribution will be made as there will be no substantial gain to deviating if this

[2]For other applications of ε-optimality see, for example, Radner et al. (1980), Baye and Morgan (2004), Barlo and Dalkiran (2009), and Milgrom (2010).

contribution is sufficiently small. All other contributions will be made in order not to forfeit the gains by others contributing in later rounds. Of course, some conditions have to be imposed on the payoff functions for this to work. However, splitting into equally-sized contributions is not very efficient in terms of the number of contributions required.

The objective of this paper is to identify the minimal number of separate contributions needed to implemented a desired total contribution. To avoid unravelling from the end, we start the construction with the last contribution. The final contribution involves a low cost for the party who makes it and high benefits for the others. 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. The fact that the last mover delivers means a lot to others who then are willing to incur costs in return earlier in the game. The efficient split is achieved by letting earlier movers invest up to the point that they are indifferent between the cost of the investment and the benefit of future investments of others. The number of rounds that are required will depend on the application. In the case of the ferryman, we can get the passenger to the other side in two rounds. 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. The ferryman anticipates the future payment and starts the journey. In other cases, the number of contributions can be large and increases as the parameter ε decreases. Typically, the size of contributions dramatically decreases towards the end. However, this is not true in general as the analysis of a three-person linear public good game shows. Throughout we assume that players are infinitely patient and that there is no cost of splitting up the contribution. In the working paper version (Martin and Schlag, 2019), we show how these two assumptions can be relaxed.

We illustrate the applicability and the simplicity of our approach through several examples. In a buyer seller game we find that payments of the buyer can be interpreted as "pay as you go". We also revisit the setting of Grossman and Hart (1986) who mitigate the hold-up problem by reassigning property rights. Under their assumptions on payoffs, we show how the hold-up problem can be eliminated entirely, even without the use of property rights.

Irreversible actions also appear in other contexts. In the "subscription game" of Admati and Perry (1991), agents care for the completion of a project in finite time, but contributions only have to be paid in case a certain threshold is reached, and hence the final investments can be self-enforcing. In contrast, investments are always and immediately costly in our setting. Some papers show that the hold-up problem vanishes if trade partners interact infinitely often and are sufficiently patient. Lockwood and Thomas (2002) consider only two-player games and cannot implement efficient outcomes. Compte and Jehiel (2004) and Che and Sákovics (2004) consider environments with bargaining and an infinite horizon. We focus on implementation once a target is fixed and importantly, our approach can be applied to realistic cases where there is a finite horizon.

Splitting up an investment was advocated in a different context by Dixit and Nalebuff (1993) in order 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 to be untrustworthy. There is no suggestion on how to split up the investment; in particular, there is no formal model. Splitting up investments also appears in job order contracting, where equally-sized parts are common practice. We show that this leads to excessively many rounds.

We proceed as follows. In Section 2, we informally present the basic principles in a two-player setting. In Section 3, we present the model. The main results for efficient implementation are in Section 4. An application to the setting of Grossman and Hart (1986) is shown in Section 5. Section 6 concludes.

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.

Consider a two-player game of complete and perfect information. The players move alternately in a given order. When a player moves, she can choose an action (investment) that is costly to herself but beneficial to others. Each time a player is making a choice she decides how much to invest. Earlier investments cannot be revoked.

The objective is to incentivize player one and player two to choose a given pair of investments (the *target*) that makes both better off compared to when neither of them makes an investment. Thus we need to find a sequence of investments that yields the desired target, and that incentivizes each player to make the desired investments. 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 or if the target is reached.

According to subgame perfect equilibrium (SPE), the last investment will not be made since it is costly to the person investing. 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, i.e., our solution concept is ε-SPE. Hence, if the last investment is not very costly to that player then she will make it. We still need to ensure that this incentivizes those moving earlier to also make their investments.

Assume that player one moves first, followed by player two. Desired investments bear a substantial cost for each player. For instance, player one as worker is supposed to make a specific observable effort, which benefits player two as employer. The employer has to pay a wage and then write a positive reference letter for player one when the relationship ends. An alternative example is the ferryman passage where both the effort of the ferryman and the fee of crossing the river are not negligible. 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. 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, her total benefits must exceed the benefits in case only player two makes the desired investment in the first round and no other investments are made. The first investment of player two is incentivized as deviation is punished by player one choosing not to invest.

To illustrate, suppose the employer first pays a good wage, then the worker exerts an observable high effort and finally the employer writes a good reference letter. The employer pays the good wage as she benefits from the effort of the worker. The worker makes the effort as the gains from the good reference letter outweigh the costs from working. The employer writes the reference letter as the time she invests in writing the letter is not very costly. In case of the river crossing we have the following scenario. The ferryman first takes the passenger across the river. Then the passenger pays the entire fee, and finally, the ferryman lets the passenger off the boat. 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, or buying a car by first receiving the official papers, then transferring the money and finally getting the key to the car.

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 this case the total investments has to be split into more parts. This is the subject of the remainder of the paper.

3 Model

For expositional clarity, we first present the model and the equilibrium concept and defer any discussion to Section 3.1 and Section 3.2.

Consider the following n-player game with alternating moves and perfect information that consists of T rounds. We denote the t-th round as the one which is t rounds from the end of the game. In each round players move in the same fixed order. Denote as player $i = \{1, 2, \ldots, n\}$ the player who is making the i-th move in each round. In each round t, each player i chooses an action $x_{i,t} \in \mathbb{R}_+$ (i.e., we consider non-reversable actions).[3] We refer to player's actions as 'investments'. The sequence of investments (*investment schedule*) is written as $(x_t)_{t=1}^T$ where $x_t := (x_{i,t})_{i=1}^n$. Each player may decide not to invest at all whenever he moves, which ends the game immediately. Each player obtains payoffs according to the total investments of each player up to that point if the game ends early, and up to the last round otherwise.

Let x^j denote the total investments of player j. We are interested in finding a value of T under which we can implement a target $\bar{x} := (\bar{x}^j)_{j=1}^n$, i.e., an allocation in which the total investments of player j equals $\bar{x}^j$. Payoffs for player i are given by $u_i\left((x^j)_{j=1}^n\right)$ where we normalize payoffs such that $u_i((0)_{j=1}^n) = 0$. Payoffs are continuous and differentiable and strictly decreasing in own investment, but increasing in the investments of the other players, formalized as follows.

Assumption 1 (Payoffs and investments)**.** *For each player i, payoffs u_i are continuously differentiable. Moreover we assume that $\partial u_i/\partial x^i < 0$ and $\partial u_i/\partial x^j \geq 0$ for $i \neq j$.*

Assumption 1 reflects the first ingredient common to hold-up problems, namely that there is a conflict between what each player wants individually versus what other players want. In order to have any hope of resolving this conflict and to implement socially desirable targets, we need a second ingredient, namely that there is path that leads to that target. Formally, we assume the following.

[3] More generally, one can allow each player a time-independent action set A_i as long as $\bar{x}_i$ (as defined below) is in the action set.

Assumption 2 (Continuous path to the target)**.** *There is a set of functions* $(h_i)_{i=1}^n$ *such that for each* $i = \{1, 2, \dots, n\}$*,* $h_i : [0,1] \mapsto \mathbb{R}_+$*,* $h_i(0) = 0$*,* $h_i(1) = \bar{x}_i$ *and both* $h_i(\tau)$ *and* $u_i(h_1(\tau), \dots, h_i(\tau), \dots, h_n(\tau))$ *are continuously increasing for* $\tau \in [0,1]$*.*

Our solution concept is ε-SPE (Mailath et al., 2005), which applies ε-Nash equilibrium to all subgames. A player i only deviates if he gains more than a given threshold ε_i. In contrast to Mailath et al. (2005), we allow for heterogeneity in ε_i and in case all player's ε_i are identical we just refer to it as $\bar{\varepsilon}$. We define $\varepsilon := (\varepsilon_i)_{i=1}^n$ where $\varepsilon_i \geq 0$ for all i and $\varepsilon_i > 0$ for at least one player i.[4] Then the concept is formally 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* $u_i(\sigma^*) \geq u_i(a_i, \sigma^*_{-i}) - \varepsilon_i$*,* $\forall a_i \in A_i$*,* $\forall i$*.*

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

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 approach when searching for outcomes that can be sustained in some ε-SPE.[5] A sequence of investments can be supported in an ε-SPE if neither player has an incentive to deviate in any round t. When deviation is punished by no more investments, we obtain that $\bar{x}$ can be supported in an ε-SPE if and only if

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

holds for all players i in all rounds t, 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 and subsequently there being no investments by players moving later.

3.1 Comments

Whether Assumption 2 holds may be difficult to verify in practice. A special case of a continuous path to the target is a *uniform continuous path*, in which $h_i(\tau) = \tau\bar{x}^i$ for each

[4]We thank an anonymous referee for suggesting to consider player-specific ε_i.

[5]Due to our irreversibility assumption, not investing in the future is the worst that players can inflict on each other and hence no subgame can lead to a lower payoff.

player i. In that case $u_i(\tau\bar{x})$ is increasing in τ on $\tau \in [0,1]$ for each player i, that is, each player benefits when all players increase their investment proportionally: $\sum_{j=1}^{n} \frac{\partial u_i}{\partial x^j} \bar{x}^j > 0$.

In Proposition 1 (proven in the Appendix), we present a simple sufficient condition that ensures existence of a continuous path for the special case of two players.

Proposition 1. *Let* $n = 2$. *If*

$$\frac{\partial u_1(x_1, x_2)}{\partial x^2} \frac{\partial u_2(x_1, x_2)}{\partial x^1} > \frac{\partial u_1(x_1, x_2)}{\partial x^1} \frac{\partial u_2(x_1, x_2)}{\partial x^2}$$

holds for all $x_1 \in (0, \bar{x}^1)$ *and all* $x_2 \in (0, \bar{x}^2)$*, then there exists a continuous path to the target* $\bar{x} = (\bar{x}^1, \bar{x}^2)$.

The assumption of Proposition 1 holds when both player's investments are socially efficient in the sense that $\frac{\partial u_i(x_1,x_2)}{\partial x^i} + \frac{\partial u_j(x_1,x_2)}{\partial x^i} > 0$ holds for $x_i \in (0, \bar{x}^i)$, $i = 1, 2$ and $j \neq i$.

Remark Both splitting up investments and ε-SPE are needed to implement any target that consists of strictly positive investments. First, suppose that we split up investments, but keep using SPE instead of ε-SPE. Then the usual unraveling argument applies. The player moving last does not make an investment since it is costly for her. Anticipating non-investment, the penultimate player also does not invest, and so on. Hence the original problem persists. Second, suppose we use ε-SPE, but do not split investments and try to implement the target in only one round. Since we typically think of ε as small relative to the total costs, the player moving last will not make an investment. Hence again the unraveling argument applies.

3.2 ε-SPE and social preferences

In this subsection we provide some motivation for our choice of ε-SPE as solution concept. We first recall the methodology underlying a Nash equilibrium (NE). In this paper we consider a commonly known strategy profile. This profile is a NE if no player has an incentive to deviate. A player deviates if and only if she can obtain strictly more when believing that she is the only one to deviate. This definition is agnostic to where a common understanding about this strategy profile comes from. The profile may be a norm for how this game is played or it can come from past interactions. SPE is a straight-forward extension of NE to extensive-form games.

In many applications, the main findings are robust to using ε-SPE instead of SPE. However, this is not the case for hold-up problems. For any strategy profile potentially implementing socially desirable outcomes in finite time, at least one player has an incentive to deviate and hence these outcomes cannot be implemented in a SPE. We will see that a slight adjustment to preferences, such that a player only deviates if gains are sufficiently large, together with our approach of splitting up investments, drastically changes the prediction. ε-SPE precisely reflects such preferences, which are very natural in many contexts.

Several reasons give rise to preferences such that a player only deviates if gains are substantial. Preferences that are close to those induced by the payoffs in the game broadly fall into two categories. The first category relates to the player herself irrespective of own payoffs. Not deviating can be based on not wishing to break a promise, the cost of contemplating the impact of this change, or the possibility of forfeiting future interactions with some of the players. Each of these can be translated into costs ε_i. More formally, consider a candidate strategy profile σ^* for implementing the desired outcome. Player i does not deviate as long as

$$u_i(\sigma^*) \geq u_i(a_i, \sigma^*_{-i}) - \varepsilon_i \tag{2}$$

holds for all $a_i \in A_i$ in every subgame, which is exactly the definition of ε-SPE.

The second category incorporates how a player feels about how own actions are perceived by others, such as costs of embarrassment, or the player does not want to be viewed by others as being greedy. The associated disutility may be captured in ε_i and leads to incentives to deviate as described in (2).

Finally, there is an even broader interpretation of ε-SPE, based on a particular form of social preferences. Consider the standard concept of an SPE and apply this to players who care a little about others, as long as no player has deviated from the equilibrium path. Then exactly the same incentive compatibility constraints emerge as in an ε-SPE. More formally, we say that a player i has *social preferences on the equilibrium path* when her preferences are represented by the utility function U_i where $U_i = u_i + \gamma \sum_{j \neq i} u_j$ for some $\gamma > 0$ on the equilibrium path and $U_i = u_i$ off the equilibrium path.

Proposition 2. *Assume that each player has social preferences on the equilibrium path with parameter $\gamma > 0$. Then there exists an $\varepsilon > 0$ such that any equilibrium that can be supported in an ε-SPE without social preferences can also be supported in a SPE with social preferences.*

Proof. For all players i, let $\phi_i = \gamma \sum_{j \neq i} u_j(\sigma^*)$ and then $\phi = \min_i \phi_i$. The assumptions on γ imply that $\phi > 0$. When players have social preference on the equilibrium path, no player deviates from a strategy profile σ^* in a SPE as long as $u_i(\sigma^*) + \phi \geq u_i(a_i, \sigma^*_{-i})$ holds for all $a_i \in A_i$, for all players i and in every subgame on the equilibrium path. This is exactly the definition of an ε-SPE for $\varepsilon = \phi$. □

Remark An alternative notion of social preferences would be that players also care about others off the equilibrium path. Then our approach still works for special payoff functions (e.g., the ferryman case with a discontinuity at the very end) but no longer works in general. For instance, in the two player linear public goods game the last investment will not be made when γ is sufficiently small. Social preferences only on the equilibrium path introduce a discontinuity in the payoff function because by deviating, the player foregoes the extra utility from other players, making the final investment incentive compatible.

4 Implementation

In this section we show how to implement any target in an ε-SPE as long as there are sufficiently many rounds. We also characterize the minimal number of rounds needed. Note that the target can be implemented when the total contribution is split up into equal parts and there are sufficiently many rounds (as formally proven in Proposition 4 in the Appendix). However, the number of rounds needed for implementation may be large in that case. We demonstrate in the following example.

Example 1 (Buyer-seller game) Consider a game between a seller (player one), making welfare-increasing investments, and a buyer (player two), making welfare-neutral payments. Payoffs are $u_1(x^1, x^2) = -x^1 + x^2$ and $u_2(x^1, x^2) = g(x^1) - x^2$ for the seller and buyer, respectively. Assumptions 1 and 2 are satisfied if g is differentiable and increasing in x, and that $g(x) - x$ is increasing up to the target investment. Consider $g(x^1) = 2\sqrt{x^1}$ that satisfies these conditions. Then it takes 50 rounds to implement the target $(1, 1)$ using constant investments when $\varepsilon = 0.01$ (i.e., 1% of the total surplus).[6]

The binding constraint is the one associated to the final investment of player two, who is the last to move in the game. He does not gain from his final investment, so the maximal investment he is willing to make does not make him more than ε worse off than not investing at all. Matters are different when it is his turn in earlier rounds, because

[6] In the last round, the buyer pays ε and the seller invests 2ε. In intermediate rounds, maximal constant payments (and hence also investments) are given by $2\sqrt{1} - 2\sqrt{1-2\varepsilon} \approx 0.02$. Thus it takes $50 = \lceil 1 + \frac{1-2\varepsilon}{2\sqrt{1}-2\sqrt{1-2\varepsilon}} \rceil \approx \lceil 49.75 \rceil$ rounds to reach a total investment of 1.

then he is incentivized by player one investing after him. A constant investment does not incorporate these dynamically changing incentives.

We now present the main result of this paper. It characterizes how many rounds are needed and shows how to determine who invests how much and when.

Proposition 3. *Fix* $(\varepsilon)_{i=1}^{n}$ *and a target* $\bar{x}$*. Assume* $\varepsilon_i \geq 0$ *for* $i < n$ *and* $\varepsilon_n > 0$*.*

i) There exists T' *such that the investment schedule* $(x_t)_{t=1}^{T}$ *implements the target* $\bar{x}$ *in an* ε*-SPE if and only if* $T \geq T'$*.*

ii) The target $\bar{x}$ *can be implemented with* $T = T'$ *by letting the incentive compatibility constraints (1) bind for all* $t < T$*.*

Proof. We first prove feasibility of implementing the target $\bar{x}$ in an ε-SPE if T is sufficiently large. Then we show that the number of rounds is minimal when incentive compatibility constraints are binding, proving both i) and ii).

In the following we show that the target $\bar{x}$ can be implemented if we split up the total contribution into sufficiently many small parts. The key is to let player n invest up to a threshold η that is close to her target and to invest only the remaining amount in the very last round. By Assumption 2, there is a continuous path to the target. We split the unit interval (i.e., the domain of τ) into smaller intervals as follows. Each subset is associated with a round s in our game, where as before $s = 1$ denotes the very last round. We denote the end point of the s-th round as τ_s. The first round starts at τ_2 and ends at $\tau_1 = 1$. Hence $\tau_s > \tau_{s+1}$. In each round s, investments for player $i < n$ are given by $x_{i,s} = h_i(\tau_s) - h_i(\tau_{s+1})$, where h_i is defined in Assumption 2. Investments of player n in rounds $s < T$ are given by $x_{n,s} = \min\{h_n(\tau_s), \eta\} - \min\{h_n(\tau_{s+1}), \eta\}$ where $\eta < \bar{x}_n$ satisfies

$$u_n\left((\bar{x}_j)_{j<n}, \eta\right) < u_n(\bar{x}) + \varepsilon_n. \tag{3}$$

Player i's period s incentive compatibility constraint (1) can be rewritten as follows:

$$\begin{aligned}
&u_i(\bar{x}) + \varepsilon_i - u_i\left((h_j(\tau_s))_{j<i}, (h_i(\tau_{s+1}))_{i\leq j<n}, \min\{\eta, h_n(\tau_{s+1})\}\right) \\
&\quad = (u_i(\bar{x}) + \varepsilon_i) - u_i\left((h_j(\tau_s))_{j<n}, \min\{\eta, h_n(\tau_s)\}\right) \\
&\qquad + u_i\left((h_j(\tau_s))_{j<n}, \min\{\eta, h_n(\tau_s)\}\right) - u_i\left((h_j(\tau_s))_{j<i}, (h_i(\tau_{s+1}))_{i\leq j<n}, \min\{\eta, h_n(\tau_{s+1})\}\right)
\end{aligned} \tag{4}$$

We first show that we can always find τ_s such that the expression in (4) is strictly positive. Define the third line in (4) as $f_i(\tau, \tau')$, hence

$$f_i(\tau, \tau') = u_i\left((h_j(\tau'))_{j<n}, \min\{\eta, h_n(\tau')\}\right) - u_i\left((h_j(\tau'))_{j<i}, (h_i(\tau))_{i\le j<n}, \min\{\eta, h_n(\tau)\}\right).$$

Since f_i can be extended to a continuous function on a compact space $[0,1]^2$, for every $\alpha > 0$ there exists $\beta_i > 0$ such that $|f(\tau, \tau')| < \alpha$ if $|\tau' - \tau| < \beta_i$.

On the other hand, we now provide arguments that the second line in (4) can be made strictly positive, independent of how τ_s is chosen. Consider a player i with $i < n$. Let $\bar{\tau}$ be such that $h_n(\bar{\tau}) = \frac{\bar{x}_n + \eta}{2}$. If $h_n(\tau) < h_n(\bar{\tau})$ then

$$u_i\left((h_j(\tau))_{j<n}, \min\{\eta, h_n(\tau)\}\right) \le u_i\left((h_j(\tau))_j\right) \le u_i\left((h_j(\bar{\tau}))_j\right) < u_i(\bar{x}).$$

Conversely, if $h_n(\tau) \ge h_n(\bar{\tau})$ then

$$u_i\left((h_j(\tau))_{j<n}, \min\{\eta, h_n(\tau)\}\right) = u_i\left((h_j(\tau))_{j<n}, \eta\right) < u_i\left((h_j(\tau))_j\right) \le u_i(\bar{x}).$$

Now consider player n. Let τ' such that $h_n(\tau') = \eta$. If $h_n(\tau) < h_n(\tau')$ then

$$u_n\left((h_j(\tau))_{j<n}, \min\{\eta, h_n(\tau)\}\right) \le u_n\left((h_j(\tau'))_{j<n}\right) < u_n(\bar{x}).$$

Conversely, if $h_n(\tau) \ge h_n(\tau')$ then

$$u_n\left((h_j(\tau))_{j<n}, \min\{\eta, h_n(\tau)\}\right) \le u_n\left((\bar{x}_j)_{j<n}, \eta\right) < u_n(\bar{x}) + \varepsilon_n.$$

Hence the second line in (4) is always strictly positive. Now define a function $g_i(\tau)$ for each player i as follows:

$$g_i(\tau) = (u_i(\bar{x}) + \varepsilon_i) - u_i\left((h_j(\tau))_{j<n}, \min\{\eta, h_n(\tau)\}\right).$$

Define $\alpha = \min_i \inf_{\tau\in[0,1]} g_i(\tau)$. The above consideration implies that $\alpha > 0$. Thus, if $\tau_s - \tau_{s-1} < \beta_i$ then

$$u_i(\bar{x}) + \varepsilon_i - u_i\left((h_j(\tau_s))_{j<i}, (h_i(\tau_{s+1}))_{i\le j<n}, \min\{\eta, h_n(\tau_{s+1})\}\right) \ge 0$$

which implies that investing is incentive compatible for player i. Thus, we can reach the target with less than $\frac{1}{\beta} + 1$ rounds in an ε-SPE where $\beta = \min_i \beta_i$.

Next we show that the implementation is shortest (in terms of number of rounds required) when we make incentive constraints binding in each round. The corresponding number of rounds is the value T' specified in the proposition. It then follows that the target $\bar{x}$ also be implemented in $T > T'$ rounds by adding zero investments in the first $T - T'$ rounds.

Consider a player i whose incentive compatibility constraint is not yet binding in some round t, whereas incentive compatibility constraints for all players are binding in the following rounds. Then the investments of all players, including player i, in the following rounds increase in the round-t investment of player i. Hence, the target $\bar{x}$ can be implemented using a smaller number of rounds.

More formally, suppose the constraint is not binding for some player i in round t, i.e., he only invests $x'_{i,t} < x_{i,t}$, but constraints are binding for all players $j < i$ in round t and also all players in all rounds $t' > t$. Binding incentive compatibility constraints for the player $i - 1$ define $x_{i-1,t}(x_{i,t})$ as the solution to

$$u_{i-1}(\bar{x}) = u_{i+1}\left(\left(\bar{x}^j - z_{j,t-1} - x_{j,t}\mathbb{1}_{\{j \geq i-1\}}\right)_{j=1}^{n}\right) - \varepsilon_i.$$

By taking implicit derivatives we obtain

$$\frac{\partial x_{i-1,t}}{\partial x_{i,t}} = -\frac{\partial u_{i-1}/\partial x_{i,t}}{\partial u_{i-1}/\partial x_{i-1,t}} > 0.$$

Analogously, it holds that $\frac{\partial x_{p,t}}{\partial x_{i,t}} > 0$ for all other players $p < i$ in round t, and players $p > i$ in round $t + 1$. Thus, other players' investments increase in $x_{i,t}$.

When player i moves again in round $t+1$, his maximal investment $x_{i,t+1}(x_{i,t})$ is given by

$$u_i(\bar{x}) = u_i\left(\left(\bar{x}^j - z_{j,t-1} - x_{j,t}(x_{i,t}) - x_{j,t+1}(x_{i,t})\mathbb{1}_{\{j \geq i\}}\right)_{j=1}^{n}\right) - \varepsilon_i.$$

By taking implicit derivatives we obtain

$$\frac{\partial x_{i,t+1}}{\partial x_{i,t}} = -\frac{\partial u_i/\partial x_{i,t} + \sum_{j<i} \partial x_{j,t}/\partial x_{i,t} + \sum_{j>i} \partial x_{j,t+1}/\partial x_{i,t}}{\partial u_i/\partial x_{i,t+1}} > -1$$

where the last inequality follows from the above observation that other players' investments increase in $x_{i,t}$. Therefore, player i's total investments across the rounds t and $t+1$ are increasing in $x_{i,t}$. Analogously, $\frac{\partial x_{j,t+1}}{\partial x_{j,t}} > 0$ for $j \neq i$. Hence, for all players it holds that total investments (weakly) increase in player i's round t investment. Therefore, the number of rounds required to implement the target is minimal when incentive compatibility constraints are binding. □

Note that the requirement that the last player is the one for whom $\varepsilon_n > 0$ holds is without loss of generality, as long as there is at least one player i with strictly positive ε_i. Simply order players such that the player moving last in each round has a strictly positive ε_i.

The intuition behind the proof of the feasibility (part (i) in Proposition 3) is as follows. Investments follow the increasing path except at the very end where the last player may deviate. In early rounds all players are willing to invest since they expect high payoffs in later rounds. Increments are sufficiently small that the payoff at the end of a round is very close to that when deviating. In the final round all players $i < n$ are waiting for the final investment of player n and hence do not deviate. Player n makes that final investment since the gains from deviating are negligible. The intuition behind the minimality result (part (ii) in Proposition 3) is as follows. Increasing an investment of some player in some round relaxes the incentive compatibility constraints of all players moving earlier in the game.

The smallest value T for implementing a given target can be of interest if there are implicit costs of splitting up a contribution. Small explicit costs can be easily included and do not change the main insights (Martin and Schlag, 2019). T is typically not very large. In the ferryman example, only two rounds are needed, regardless how small ε is. For $g(x^1) = 2\sqrt{x^1}$ in Example 1, 18 rounds suffice to implement the target $(1, 1)$ when $\bar{\varepsilon} = 0.01$. Recall that 50 rounds are needed under constant investments).

Our approach yields a particularly simple solution, as illustrated by the following example.

Example 1 (continued) The optimal investment schedule is given by $x_{2,t} = x_{1,t} = g\left(\sum_{s=t-1}^{T} x_{1,s}\right) - g\left(\sum_{s=t}^{T} x_{1,s}\right)$ for t where $2 \leq t \leq T-1$, $x_{1,1} = 2\bar{\varepsilon}$ and $x_{2,1} = \bar{\varepsilon}$, and at the beginning by whatever is left to reach the target, that is, $x_{1,T} = \bar{x}^1 - \sum_{i=1}^{T-1} x_{1,i}$ and $x_{2,T} = \bar{x}^2 - \sum_{i=1}^{T-1} x_{2,i}$. Thus, the buyer's payment $x_{2,t}$ exactly compensates the seller for his investment $x_{1,t}$ ('pay as you go'). The seller invests in round t exactly the incremental utility for the buyer generated by the investment in the next round $t-1$. As we show in the working paper version (Martin and Schlag, 2019), a similar structure emerges in settings where payoff function exhibit constant returns to scale.

The above considerations may suggest that investments are decreasing closer to the deadline. This is true for most of the numerical examples we consider in Martin and

Schlag (2019). We now present an example where this intuition fails.

Example 2 Consider a three-player symmetric linear public goods game, 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 $1/3 < \alpha < 1$. Hence $\partial u_i/\partial x_i = \alpha - 1$ and $\partial u_i/\partial x_j = \alpha$. Thus this game gives satisfies our assumptions. The target of $(1,1,1)$ maximizes the sum of the players' payoffs and can be reached with a uniform continuous path. According to our approach, investments in round 1 are given by $x_{1,1} = \frac{\bar{\varepsilon}}{(1-\alpha)^3}$, $x_{2,1} = \frac{\bar{\varepsilon}}{(1-\alpha)^2}$, $x_{3,1} = \frac{\bar{\varepsilon}}{1-\alpha}$, and in round 2 by $x_{1,2} = \frac{(3-2\alpha)\alpha^2\bar{\varepsilon}}{(1-\alpha)^6}$, $x_{2,2} = \frac{(1+\alpha-\alpha^2)\alpha\bar{\varepsilon}}{(1-\alpha)^5}$ and $x_{3,2} = \frac{(2-\alpha)\alpha\bar{\varepsilon}}{(1-\alpha)^4}$. It is easily verified that investments are decreasing in time if and only if $\alpha > 0.3473$. When α is below this threshold, we have $x_{1,1} > x_{2,1}$ and hence investments are not decreasing. When player one moves in round one (the last round), he invests much more than the other players moving after him because he has more to lose given the benefits by later investments by other players. When player one moves in round two, his incentives are dampened by his own large future investment. This would not matter when the round two investments of the other players are high. However, when α is small, the others invest too little in round two. As a consequence player one reduces his round two investment relative to the round one investment.

5 Grossman and Hart (1986)

In this section, we apply our insights to the hold-up problem described in Grossman and Hart (1986), who mitigate the hold-up problem by introducing property rights. We consider the leading example in this literature as presented in Hart (1995, Chapter 2). We show how to solve the hold-up problem by appropriately splitting up investments.

There are two players who can make an investment and exert effort prior to negotiating terms of trade. First, player one can make an investment i to increase revenue $R(i)$ and simultaneously player two can exert effort to decrease costs $C(e)$. Then the players negotiate the transaction price p. Alternatively they can trade on the open market at a market price $\bar{p}$ and with revenue $r(i)$ and production costs $c(e)$, respectively. Both investment and effort are relationship-specific and non-verifiable and therefore there is a hold-up problem.

If the two players trade at price p, payoffs are given by $u_1 = R(i) - p - i$ and $u_2 = p - C(e) - e$, for player one and two, respectively. Trading on the open market results in payoffs $u_1 = r(i) - \bar{p} - i$ and $u_2 = \bar{p} - c(e) - e$. The following conditions are assumed to hold on the relevant domain: $R' > r' > 0$, $C' < c' < 0$, $R'' < 0$, $C'' > 0$. Given these conditions, players prefer trading among themselves to trading on the market. Prices are determined by Nash bargaining. Thus the market price p is given by $p(i,e) = \bar{p} + \frac{R(i)-r(i)+C(e)-c(e)}{2}$ and hence payoffs are given by $u_1(i,e) = -\bar{p} + \frac{r(i)+R(i)-C(e)+c(e)}{2} - i$ and $u_2(i,e) = \bar{p} - \frac{r(i)-R(i)+C(e)+c(e)}{2} - e$.

Grossman and Hart (1986) show that the efficient outcome cannot be implemented, but property rights mitigate the problem. We show that all the conditions required to make our approach work are satisfied, and hence the efficient outcome can be implemented by splitting up investments.

The target is the social optimum (i^*, e^*), i.e., the investment and effort pair that maximizes joint surplus. This has to satisfy $R'(i) = 1$ and $C'(e) = -1$. The Nash equilibrium (i^N, e^N) of the game is given as solution to $\frac{r'(i)+R'(i)}{2} = 1$ and $\frac{C'(e)+c'(e)}{2} = -1$. Given our assumptions, $(i^N, e^N) < (i^*, e^*)$ (i.e., there is under-investment relative to the social optimum).

Even without splitting up investments, Nash equilibrium levels can be implemented by definition of a Nash equilibrium. In the following, we only consider the subgame that emerges once Nash equilibrium investments were already made.[7] Beyond Nash equilibrium levels, this setting satisfies all our assumptions. $\partial u_1/\partial e > 0$, $\partial u_2/\partial i > 0$ and moreover for $i \in (i^N, i^*)$, $\partial u_1/\partial i < 0$ and for $e \in (e^N, e^*)$, $u_2/\partial e < 0$. Thus, Assumption 1 holds. By definition of the socially efficient target, investments of either player are socially optimal up to that point (i.e. $\frac{\partial u_1}{\partial i} > -\frac{\partial u_2}{\partial i}$ and $\frac{\partial u_1}{\partial e} > -\frac{\partial u_2}{\partial e}$) and hence

$$\frac{\partial u_1}{\partial e}\frac{\partial u_2}{\partial i} > \left(-\frac{\partial u_1}{\partial i}\right)\left(-\frac{\partial u_2}{\partial e}\right) = \frac{\partial u_1}{\partial i}\frac{\partial u_2}{\partial e}.$$

This is exactly the condition of Proposition 1, and therefore a continuous path to the target exists. Hence all our conditions are satisfied, and by Proposition 3, the target can be implemented in an ε-SPE.

Example 3 Let $R(x) = -C(x) = 2\sqrt{x}$, $r(x) = -c(x) = \sqrt{x}$ under non-integration, $r_1 = 3/2\sqrt{x}$ under type-1 integration where firm 1 obtains ownership over all assets, and arbitrary $\bar{p}$. This is the most favorable case for integration since the incentives for player 2 do not change, that is, c remains unaffected. Then $i^* = e^* = 1$ and $i^N = e^N = 9/16$. Total welfare in the Nash equilibrium (non-integration) is given by $W^N = u_1^N + u_2^N = 1.875$ and

[7] This simplifies the exposition. Our results generalize to a setting in which players initially benefit from their own actions.

under type-1 integration $W^I = 1.92$. The efficient outcome yields welfare $W^* = 2$. Thus, property rights eliminate only roughly one-third of the inefficiency of the Nash outcome, whereas our approach eliminates it entirely.

6 Conclusion

In this paper we introduce a general and novel method for implementing socially desirable outcomes within finite horizons without any use of institutions. It is applicable if the entire problem can be split up into several smaller problems that can then be executed in sequence, such that all information about previous interactions is available at each choice. It is implemented with ε-SPE as solution concept. We demonstrate this method in environments in which inaction (no investment) is preferred to any other choice and own choices benefit others. In our exposition we focus on the case where players are infinitely patient and where there are no costs of splitting up decisions. In the working paper version (Martin and Schlag, 2019) we show that all results extend generically to the setting where players are sufficiently patient and the costs of splitting up investments are sufficiently small.

Once the two ingredients ε-SPE and splitting up transactions are identified, it is easy to see that the hold-up-problem can be solved. However, this paper goes substantially further by showing how it can be solved with the minimal number of rounds and thereby provides novel insights.

Our paper contributes on three dimensions: theory, description and design. On the theoretical or normative side, our paper initiates by adding two new building blocks to the finite horizon hold-up, ε-SPE and splitting up decisions. In particular, we object to a SPE as a methodology that rules out outcomes, simply because some player could have made very small gains. We provide many reasons why such gains will not be taken. The results hold for any strictly positive ε. As a noteworthy theoretical insight, we find that an arbitrarily small possibility of a continuing relationship suffices to sustain cooperation. This stands in contrast to our lessons from repeated games whereby future stakes and interaction probabilities will need to be sufficiently large if one wishes to enforce cooperation. From the theoretical perspective, our contribution can be identified by finding a clean characterization of the optimal solution. As a methodological contribution, the simplicity of the analysis and the characterization makes examples easy to calculate. Insights are accessible for wider audiences; insights that are easily used in teaching strategic decision making and game theory. Open topics left for future research include understanding how investments change over time.

From the descriptive viewpoint we clarify 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 to incentivize the earlier, more substantial investments.

Our paper also initiates a more applied research agenda on transaction design. 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 covered either of these topics as it deals with a single dimensional investment. For example, a landlord may ask for a small present, such as a refrigerator magnet, after the lease has terminated. This creates a future benefit that he cares about and that can be provided at a low cost by the tenant. 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 for using 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.

We end our exposition with some observations that emerge from this paper. There seems to be a life cycle when attempting to create surplus with costly interactions. One might refer to reverse gradualism. Early investments are large as they are incentivized by later ones. As the relationship carries on, investments become smaller to reduce gains of deviation. At the end of the project some small rewards are required in order to incentivize the last choices that are needed in a long chain to incentivize the earlier ones. Incentives are close knit and interrelated at each of the decision points. Small items at the end of a relationship can have a significant impact, such as handing over the key to the house. Large investments no longer have to be incentivized if they are followed by others, as all of the focus lies on the ability to incentivize that last investment. This can be done by creating a future, such as opening the possibility of another interaction, introducing a handshake to create emotional attachment, introducing alumni organizations, or simply by sending Christmas cards.

Appendix A: Proofs

Proof of Proposition 1. We implicitly define x_1 and x_2 as functions of $y \in [0, u_1(\bar{x})]$ as follows:

$$u_1(x_1(y), x_2(y)) - u_1(0,0) - y = 0 \tag{5}$$

$$u_2(x_1(y), x_2(y)) - u_2(0,0) - \gamma y = 0 \tag{6}$$

where $\gamma = \frac{u_2(\bar{x}) - u_2(0,0)}{u_1(\bar{x}) - u_1(0,0)} > 0$. When solutions $x_1(y)$ and $x_2(y)$ exist, then they are, by construction, a continuous path to the target.

Define the matrix DF as

$$DF(x_1, x_2) = \begin{pmatrix} \frac{\partial u_1(x_1,x_2)}{\partial x^1} & \frac{\partial u_1(x_1,x_2)}{\partial x^2} \\ \frac{\partial u_2(x_1,x_2)}{\partial x^1} & \frac{\partial u_2(x_1,x_2)}{\partial x^2} \end{pmatrix}.$$

By continuity of the utility functions and the Implicit Function Theorem, $x_1(y)$ and $x_2(y)$ exist and are solutions to (5) and (6) when the determinant $\det(DF(x_1, x_2))$ is non-zero in its entire domain:

$$\det(DF(x_1, x_2)) = \frac{\partial u_1(x_1, x_2)}{\partial x^1} \frac{\partial u_2(x_1, x_2)}{\partial x^2} - \frac{\partial u_1(x_1, x_2)}{\partial x^2} \frac{\partial u_2(x_1, x_2)}{\partial x^1}$$

which is strictly negative by assumption.

Our assumption of an increasing path to the target also requires increasing investments along the path, hence we also need that x_1 and x_2 are (weakly) increasing functions of y. The derivatives are given by

$$\begin{pmatrix} x_1'(y) \\ x_2'(y) \end{pmatrix} = -(DF(x_1, x_2))^{-1} \begin{pmatrix} -1 \\ -\gamma \end{pmatrix}$$

where

$$(DF(x_1, x_2))^{-1} = \frac{1}{\det(DF(x_1, x_2))} \begin{pmatrix} \frac{\partial u_2(x_1,x_2)}{\partial x^2} & -\frac{\partial u_1(x_1,x_2)}{\partial x^2} \\ -\frac{\partial u_2(x_1,x_2)}{\partial x^1} & \frac{\partial u_1(x_1,x_2)}{\partial x^1} \end{pmatrix}.$$

Since the determinant is negative, $x_1'(y)$ is positive if

$$\frac{\partial u_2(x_1, x_2)}{\partial x^2} - \gamma \frac{\partial u_1(x_1, x_2)}{\partial x^2} \leq 0$$

which holds by Assumption 1. $x_2'(y)$ is positive if

$$-\frac{\partial u_2(x_1, x_2)}{\partial x^1} + \gamma \frac{\partial u_1(x_1, x_2)}{\partial x^1} \leq 0$$

which also holds by Assumption 1. Thus, there exists a continuous path to the target. □

Proposition 4. *Fix $\bar{\varepsilon} > 0$ and a target $\bar{x}$ and suppose there is a uniform continuous path to the target. Then there exist T' and constants $(c_i)_{i=1}^n$ where $c_i > 0$ for each player i and T such that using constant investments $x_{i,t} = c_i$ in each round $t < T$, the target can be implemented in an ε-SPE if and only if $T \geq T'$.*

Proof. By the uniform continuous path assumption, we have that $h_i(\tau) = \tau \bar{x}^i$ for each player i, and hence $u_i(\tau \bar{x})$ is increasing in τ on $\tau \in [0, 1]$ for each player i.

We first find a fraction $\phi \in (0, 1)$ such that it is incentive compatible for each player i to invest a fraction ϕ of his target $\bar{x}^i$ in the last period while everyone else doesn't invest. Define ϕ_i as the solution to $u_i(\bar{x}) = u_i\left(\left(\bar{x}^j - \phi_i \bar{x}^j \mathbb{1}_{\{j=i\}}\right)_{j=1}^n\right) - \bar{\varepsilon}$ and then ϕ as $\phi := \min_{i=1}^n \phi_i$. Now suppose each player invests a constant $c_i = \phi \bar{x}^i$ in each period. This is incentive compatible for each player i in each round t since it holds that

$$u_i(\bar{x}) \geq u_i(t\phi\bar{x}) \geq u_i\left(\left(\bar{x}^j - (t-1)\phi\bar{x}^j - \phi\bar{x}^j \mathbb{1}_{\{j \geq i\}}\right)_{j=1}^n\right) - \bar{\varepsilon}$$

where the first inequality follows from the uniform continuous path assumption and the second inequality follows the definition of ϕ. Hence this schedule implements the target in an ε-SPE if and only $T \geq T' := \lceil 1/\phi \rceil$.

□

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.

Grossman, S. J. and Hart, O. D. (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy*, 94(4):691–719.

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.

Martin, S. and Schlag, K. H. (2019). Split it up to create incentives: Investment, public goods and crossing the river.

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.