// NO.25-046 | 09/2025

DISCUSSION PAPER

// BERNHARD GANGLMAIR, JULIAN KLIX, AND DONGSOO SHIN

Hybrid Contracting in Repeated Interactions

Hybrid Contracting in Repeated Interactions*

Bernhard Ganglmair† Julian Klix‡ Dongsoo Shin§

August 29, 2025

Abstract

Many business relationships begin with informal interactions and later transition to formal contracts. Using a repeated-games model with a finite horizon, we show that this hybrid-contracting approach can both prolong cooperation (intensive margin) and enable it across a broader range of settings (extensive margin). We model the contract as a "smooth-landing contract" that limits actions only near the end of the relationship. We show that this flexible design supports early cooperation and outperforms rigid contracts. Our findings are robust to changes in contracting costs and timing, with optimal contract length balancing profitability and implementability.

Keywords: cooperation, hybrid contracting, relational contracts, repeated games, strategic alliances

JEL Codes: C73, D86, K12, L14

---

*We thank conference and seminar participants at EARIE, the Stony Brook Game Theory Conference, and Jacopo Gambato and Nicolas Schutz for helpful comments and suggestions. Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project B02) is gratefully acknowledged. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

†University of Mannheim and ZEW Mannheim, Germany. E-mail: b.ganglmair@gmail.com. Phone: +49-621-1235-304 (corresponding author)

‡University of Mannheim, Germany. E-mail: julian.j.klix@gmail.com

§Santa Clara University, Santa Clara, USA. E-mail: dshin@scu.edu

1 Introduction

The role of cooperation and relational dynamics is widely recognized in economics and management (e.g., Baker et al., 2002). Practitioners frequently point out that being better at using contracts starts with good strategies to foster trust and cooperation—that is, building trust and using contracts are complementary to each other. Along these lines, Ohmae (1989) reports that many business alliances (such as strategic alliances, research joint ventures, or other collaborative agreements) depend on softer aspects of relationships, noting that few rely on rigidly binding agreements from the beginning. Instead, he emphasizes the importance of loose arrangements in the early stages only to contractually solidify business relationships later. Similarly, firms regularly offer small and standardized contracts to vet their suppliers before moving to larger and more complex orders (Moore et al., 2002; Liker and Choi, 2004; Bernstein and Peterson, 2022, fn. 209).

In this paper, we study how firms can use a progression from relational contracts at the early stages of a relationship to formal contracting (and commitment) at later stages to ensure mutual cooperation throughout the entire relationship. To model such hybrid contracting, we utilize a "finitely" repeated game (e.g., Benoit and Krishna, 1985; Fudenberg and Tirole, 1991), building on (partial-)cooperation results where firms cooperate at all early stages but cooperation breaks down toward the end. Unique to our modeling approach is that firms can commit to a formal contract, referred to as a *smooth-landing contract*, that restricts the strategy space at those later stages.

We demonstrate that a smooth-landing contract can both extend the duration of a cooperative business relationship (intensive margin) and expand the set of environments in which cooperation can be achieved in the first place (extensive margin). Our analysis reveals that, by restricting the action space, a contract can retain and complement the early-stage relational dynamics. It thus extends cooperation throughout the entire business relationship more cost-effectively than a benchmark full-commitment contract that prescribes actions at all stages of the relationship (and does not make use of the early-stage relational dynamics). Moreover, an abridged full-commitment contract that prescribes the cooperative action only at the later stages (when the relational dynamics otherwise fail), regardless of past actions, ends up crowding out the early-stage cooperative behavior under the relational dynamics.

Our model represents a phenomenon frequently encountered in the context of collaborative agreements and alliances between competitors, where relationships are initially kept loose but solidified later. Firms seek out collaborative agreements to improve their market positions or reduce rivalry (Park and Ungson, 2001; Schilling, 2019, ch. 8), and such alliances are on the rise (Gulati et al., 1994; Park and Russo, 1996). However, a striking 60–70% of alliances fail (Hughes and Weiss, 2007), often "when excessive rivalry eclipses cooperative tendencies" (Park and Ungson, 2001, p. 38). More commitment can overcome the loose structure of an alliance and address potential opportunism by alliance partners. An example is an R&D joint venture, a particular type of strategic alliance characterized by significant structure and commitment that goes beyond that of a simple contractual or even non-contractual alliance (Gulati and Singh, 1998), often resulting in a separate entity (Schilling, 2019, p. 167).[1] In our stylized model of strategic alliances, a smooth-landing contract plays the role of a contractual alliance or R&D joint venture, implemented to overcome the opportunism at the later stages of the relationship.

Serving a strategic purpose similar to the contractualization of alliances, managerial contract provisions can facilitate cooperative contracting relationships.[2] Moore et al. (2002) and Liker and Choi (2004) describe how firms frequently vet their suppliers with small and standardized contracts before moving to larger and more complex orders to establish a pattern of cooperative behavior. Such managerial governance techniques help maintain these cooperative relationships, also mitigating the risk that an inadvertent non-performance by one side is not interpreted as a defection (met with a defection in return, ultimately leading to a breakdown of the relationship) by the other (Bernstein and Peterson, 2022, pp. 221–223).

For our benchmark results, we model an ongoing relationship between two parties as a symmetric, finitely repeated game that can generate partial cooperation in the early stages. To that end, our stage game, inspired by a prisoner's dilemma game, features three actions. This extension of the classical two-action prisoner's dilemma allows our model to be as concise as possible while avoiding the complete backward unraveling—and thus complete lack of cooperation—that would be inherent to such two-action stage games. In our setup, a *grim trigger* strategy (that means, unforgivingly punishing deviation) forms the baseline that generates a partial-cooperative equilibrium where players cooperate at all early stages as long as they do not discount the future too much (similar to Benoit and Krishna (1985)), and a non-cooperative equilibrium otherwise.[3]

We then introduce contracts in the form of commitment devices that allow players to commit to playing (or avoiding) specific actions. The goal of such agreements is to extend the stages in which cooperation takes place (*intensive margin*) or the parameters in which cooperation can occur entirely (*extensive margin*). What we refer to as *smooth-landing contracts* is one particular class of such incomplete contracts, leaving the players' hands free at the early stages and only coming into effect at the late stages of the game. This way, smooth-landing contracts can prevent the backward unraveling that precludes late-

[1]Formalizing alliances as joint ventures may also help firms avoid antitrust scrutiny (Strachan, 2019).

[2]Cooperative contracting relationships are "relationships where shirking is minimized, relationship-specific investments are adequately bonded, and opportunistic behavior is adequately controlled" (Bernstein, 2016, p. 576).

[3]In fact, the limited folk theorem using trigger strategies proposed by Benoit and Krishna (1985) (see also Fudenberg and Tirole (1991)), though a particularly tractable one, is not the only approach to obtain games with partial cooperation behavior. For another approach using incomplete information, see Kreps et al. (1982).

stage cooperation without eliminating the incentives used by the grim trigger strategies that generate cooperation already in the baseline.

Smooth-landing contracts can broaden cooperation at the *extensive margin* (i.e., when there has been no cooperation before) as well as the *intensive margin* (i.e., increasing the periods of cooperation) and are less costly than the full-commitment alternative because of their shorter length and later enforcement. Hence, we find that whenever they are profitable (that means, better than the commitment-free baseline with partial or no cooperation), smooth-landing contracts are also optimal within our framework (as they cost-dominate, for instance, full-commitment contracts).

In extensions of the baseline model, we demonstrate that our results are robust to incorporating more complex contracting costs (allowing for ex-ante drafting costs in addition to ex-post enforcement costs) and endogenizing the timing of contract negotiation. We further show that the optimal length of smooth-landing contracts balances profitability (and contracting costs) with the range of implementability of the cooperative outcome.

Related Literature. Our analysis contributes to the literature on contracting and strategic alliances in several ways. First, our results provide a rationale for the common practice of maintaining loose relationships (potentially relying on relational dynamics) initially, before solidifying them with formal contracts and commitments. For alliance management practitioners, our results highlight the potential of formal contracting to complement relational dynamics cost-effectively by filling gaps rather than replacing relational dynamics with formal ones.[4] Our results also stress the importance of designing flexible smooth-landing contracts to avoid unintentionally crowding out the very relational dynamics they are meant to complement. The absence of smooth-landing contracts or the use of rigid contracts results in alliance failure in our model, adding to a list of factors identified in the literature (e.g., Park and Ungson, 2001) that are responsible for the high failure rate of business alliances (Parkhe, 1993; Hughes and Weiss, 2007). Our model rationalizes contractual solutions designed to overcome alliance failures when initially loose alliances are successful.

Second, our baseline setup, which involves cooperation in a finitely repeated game, is related to work in the literature on relational contracts. In finitely repeated games, cooperation is typically not attainable because of the underlying end-of-game properties. Yet, experimental evidence shows that cooperation is possible in such settings (Andreoni and Miller, 1993). Kreps et al. (1982) establish cooperation in a finitely repeated prisoner's dilemma by introducing reputation and asymmetric information. We adopt the approach outlined in Benoit and Krishna (1985) to achieve some cooperation in our baseline model and then incorporate formal commitment into the picture. Studies by Baker

[4]For empirical evidence of the complementary use of relational and formal contracting, see Ryall and Sampson (2009).

et al. (2002), Poppo and Zenger (2002), and Gibbons and Henderson (2012) are similar in spirit. Their main finding is that formal contracts complement, rather than replace, relational (informal) contracts. Unlike in our paper, however, the action set in their papers comprises both formally contracted components and left-to-repeated-relationship components. Although the studies also show that infusing formal contracts with relational elements enhances cooperation, the main point differs from ours. Our study focuses on the dynamics of a *shift* or *progression* in the contracting structure—when it is better to choose contracts that complement the relational dynamics instead of replacing them or crowding them out altogether—whereas their work focuses on the *mixture* of the structure in a stage game that is played repeatedly.[5]

Third, our stylized model of informal and formal contracting relates to the literature on incomplete contracts, particularly in the context of the question of when and why parties choose to write more or less complete contracts. Contractual incompleteness is often associated with transaction costs (Williamson, 1985, 1989) that can arise both ex ante and ex post. Ex-ante transaction costs are the costs associated with the drafting of the contract ("search costs" as in Klein (2002) or Tirole (2009) and "ink costs" as in Dye (1985), Anderlini and Felli (1994), Melumad et al. (1997), or Battigalli and Maggi (2002, 2008)). Ex-post transaction costs are the costs of enforcing and implementing contracts (Kaplow, 1995; Schwartz and Watson, 2004).

In our main model, we assume fixed ex-post enforcement costs (but no ex-ante drafting costs) that do not depend on the characteristics or duration of the contract. Contracting parties incur these costs only once, and because of time discounting, they prefer contractual commitments in later periods. This adds a *temporal* dimension to optimal contractual incompleteness—enforcement costs of longer and earlier contracts are effectively higher—thereby introducing a trade-off between wider implementability and higher costs. Instead of writing a full (or complete) contract, parties can agree on a more cost-effective incomplete contract now and fill in the gaps later. In our setting, the relational part is akin to the gap in an incomplete contract, and the smooth-landing contract fills that gap later. In an extension, we demonstrate that this logic remains intact when allowing for both ex-post (enforcement) and ex-ante (drafting) costs.

Relatedly, Battigalli and Maggi (2002) argue that there may be trade-offs between additional detail and overly rigid contracts, where contractual contingencies may be insufficiently customized to the contracting parties. They show that an optimal contract may leave some degree of discretion to the parties—that means the contract is incomplete.[6] In line with these results, we demonstrate that flexible smooth-landing contracts with a larger allowable action space outperform an abridged (and rigid) full-commitment

[5] Also, Baker et al. (2002) and Gibbons and Henderson (2012) employ a model with infinite horizon, whereas we adopt a finitely repeated setting. Poppo and Zenger (2002) demonstrate complementarity between the contractual structures empirically.

[6] Basov (2016, pp. 62–64) provides a simple example to illustrate this point.

contract that effectively prescribes cooperative actions. This is because more rigid contracts are unable to maintain the cooperative relational dynamics of the baseline (and hybrid contracting).

The remainder of the paper is structured as follows: In Section 2, we analyze the baseline model and characterize the benchmark equilibria without contracts. In Section 3, we formally introduce hybrid contracting and derive the parameter space in which smooth-landing contracts can increase cooperation relative to the benchmark equilibria. In Section 4, we characterize the implementability of the cooperative outcome under contractual solutions, determine their profitability, and derive the parties' preferred solution that maximizes joint profits. In Section 5, we extend the baseline model to allow for more flexible timing and contracting costs. In Section 6, we conclude.

2 Baseline Model of Repeated Interaction

2.1 Setup

Let there be two players interacting in a finitely repeated normal-form game. As is well known, early cooperation between the players requires multiple Nash equilibria in a stage game (Benoit and Krishna, 1985). For simplicity, we focus on a repeated game with three possible actions, resulting in an *augmented prisoner's dilemma* as its stage game. This approach allows using a grim trigger strategy of Nash reversion to sustain the cooperative outcome in early periods with a collapse of cooperation in the last period (we will refer to this equilibrium as a *partial-cooperative* equilibrium).

The stage game consists of two players $i = 1, 2$ deciding whether to work (w), loaf (l), or shirk (s).[7] Each action a_i in $A_i = \{w, l, s\}$ is associated with an effort level $e(a_i)$, where $e(w) = 2$, $e(l) = \tilde{e}$, and $e(s) = 0$. The payoff in each stage-game outcome is then given by the function

$$\pi_i(a_i, a_j) = [e(a_i) + e(a_j)]\,\pi - e(a_i)c - \mathcal{I}(e(a_i) < e(a_j)) \cdot m, \tag{1}$$

where π is the utility of joint effort, c is the cost of effort, m is the moral utility cost of free-riding, and $\mathcal{I}(\cdot)$ is an indicator function equal to one if $e(a_i) < e(a_j)$, zero otherwise.[8] To simplify the payoff structure, we normalize the payoffs for the (w, w) outcome to 1 by setting π and c such that $\pi_i(w, w) = 4\pi - 2c = 1$. We summarize the symmetric payoff matrix of the stage game in Table 1.

[7]In this game w dominates s, whereas the intermediate action l is neither dominant nor dominated. Thus, we refer to this structure as an augmented prisoner's dilemma.

[8]The existence of a moral or conscience cost that reduces a player's inclination to free ride is well established in the public economics literature (e.g. Cubitt et al., 2011).

Table 1: Normal-Form Representation of the Stage Game

		Player 2		
		s (shirk)	l (loaf)	w (work)
Player 1	s (shirk)	$\begin{pmatrix} 0 \\ 0 \end{pmatrix}$	$\begin{pmatrix} \tilde{e}\pi - m \\ \tilde{e}(\pi - c) \end{pmatrix}$	$\begin{pmatrix} 2\pi - m \\ 2(\pi - c) \end{pmatrix}$
	l (loaf)	$\begin{pmatrix} \tilde{e}(\pi - c) \\ \tilde{e}\pi - m \end{pmatrix}$	$\begin{pmatrix} \tilde{e}/2 \\ \tilde{e}/2 \end{pmatrix}$	$\begin{pmatrix} (2+\tilde{e})\pi - \tilde{e}c - m \\ (2+\tilde{e})\pi - 2c \end{pmatrix}$
	w (work)	$\begin{pmatrix} 2(\pi - c) \\ 2\pi - m \end{pmatrix}$	$\begin{pmatrix} (2+\tilde{e})\pi - 2c \\ (2+\tilde{e})\pi - \tilde{e}c - m \end{pmatrix}$	$\begin{pmatrix} 1 \\ 1 \end{pmatrix}$

To ensure that (s, s) and (l, l) form the only two pure-strategy Nash equilibria in the stage game, we impose the following conditions:[9]

Condition 1. *Let the set of parameters* $(\pi, c, m, \tilde{e})$ *be such that (a)* $\pi > \frac{1}{2}$*; (b)* $c = 2\pi - \frac{1}{2}$*; (c)* $0 < m < 2\pi - 1$*; and (d)* $0 < \tilde{e} < \min\left\{\frac{m}{\pi - 1/2}, \frac{m+1/2}{\pi}\right\}$.

For better tractability, let $\pi^D = \pi_i(s, w) = 2\pi - m$ denote the maximal deviation payoffs and $\pi^L = \pi_i(l, l) = \tilde{e}/2$ denote the payoffs of the (l, l) stage-game Nash equilibrium. Moreover, Condition 1(b) normalizes the payoffs for (w, w) to 1. Note that $\pi^D > 1$ by Condition 1(c) and $\pi^L < 1$ by Condition 1(c) and (d). The augmented prisoner's dilemma stage game is repeated T times (for a total of $T + 1$ rounds) with a standard discounted total utility $\Pi_i = \sum \delta^t \pi_i(a_{i,t}, a_{j,t})$. We make the standard assumptions of rational players and common knowledge of past actions.

2.2 Equilibria

Because the augmented prisoner's dilemma stage game has two pure-strategy Nash equilibria, consider the following strategy to form a *partial-cooperative equilibrium*:

Baseline Strategy. *If all parties have cooperated without failure, continue to cooperate and play* w*. If a party has deviated at any point, play* s *(as a punishment) for the remainder of the game. In the last period, respond to uninterrupted cooperation by playing* l*.*

[9] These assumptions also guarantee the existence of only one mixed-strategy Nash equilibrium in the stage game. For better tractability, we focus our analysis solely on pure-strategy Nash equilibria. Allowing for mixed strategies in our analysis, however, does not qualitatively alter the results. For a discussion of the effects of mixed-strategy equilibria, see Appendix A.3. Alternatively, the existence does not affect the outcomes under the modified condition for *(d)*, that is $\frac{m}{\pi} < \tilde{e} < \min\left\{\frac{m}{\pi - 1/2}, \frac{m+1/2}{\pi}\right\}$.

Formally, this is saying that every party plays the grim trigger strategy

$$a_{i,t}^{GT} = \begin{cases} s & \textit{if } \exists t' < t : a_{j,t'} \neq w; \\ l & \textit{if } t = T \wedge \forall t' < t : a_{j,t'} = w; \\ w & \textit{otherwise.} \end{cases} \tag{2}$$

Note that in a repeated game, playing anything other than a stage-game Nash equilibrium in the last stage would allow for a profitable deviation without affecting previous stage-payoffs. The equilibrium action profile in the last period T must, therefore, be (l, l) (rather than (w, w)) with payoffs π^L even if both players cooperated up until the last period.

A player finds it profitable to deviate as late as possible if

$$\delta > \frac{\pi^D - 1}{\pi^D} =: \bar{\delta}^D. \tag{3}$$

Given a discount factor satisfying this condition, a deviation in $T - 1$ (and any earlier periods) is unprofitable whenever

$$\delta \geq \frac{\pi^D - 1}{\pi^L} =: \bar{\delta}^{BL}. \tag{4}$$

This inequality yields the equilibrium condition for a partial-cooperative equilibrium (with cooperation in all but the last period), summarized in the following lemma:

Lemma 1. *A partial-cooperative equilibrium, yielding the outcome* $(w, w), \ldots, (w, w), (l, l)$*, exists in the baseline model whenever* $\delta \geq \frac{\pi^D - 1}{\pi^L} =: \bar{\delta}^{BL}$.

The idea for the proof of Lemma 1 is straightforward. Because the outcome in T is the Nash equilibrium (l, l), we are left to show that the player's deviation incentives are strongest in $T - 1$ to obtain the critical threshold in the lemma. If the player cooperates in that penultimate period, she cooperates in all earlier periods. As the decision moves up to earlier periods, the duration of punishment and, therefore, the opportunity costs from a deviation increase, while the benefits from deviation stay constant. As a consequence, a deviation becomes less profitable for smaller values of t (in earlier periods).[10]

In addition to this partial-cooperative equilibrium, it is also always optimal (in the sense of a *subgame-perfect Nash equilibrium*) to play a Nash equilibrium in every period. That means that every action profile $a_1^* = a_2^* \in \{s, l\}^{T+1}$ also constitutes a subgame-perfect Nash equilibrium of the augmented prisoner's dilemma, of which $a_i^* = (l, \ldots, l)$ is payoff-dominant, forming the *non-cooperative equilibrium* characterized in Lemma 2.

[10]Note that the dynamics are different in an infinitely repeated game where the opportunity costs from a deviation in a period t are independent of the actual t.

Lemma 2. *A non-cooperative equilibrium, yielding the outcome $(l,l),\ldots,(l,l)$, exists in the baseline model for any δ.*

This non-cooperative equilibrium is of particular relevance in the case when a partial-cooperative equilibrium cannot be sustained, that is, whenever $\delta < \bar{\delta}^{BL}$.

3 Commitment Through Contracting

While partial cooperation can be sustained in the baseline model, adding the ability to write contracts to the framework will give the players additional devices to improve cooperation—both *intensively* (increasing the number of periods with cooperation) and *extensively* (extending the region of the parameter space in which cooperation is possible). Assuming that actions are contractible, a contract allows the players to commit to a certain subset of actions available to play for a specified subset of periods. The allowable actions from this contract may depend on earlier actions taken.

For our main results, we assume that any contract is entered just before the first interaction (in $t = 0$).[11] At the outset of the game, the players contract to specify a subset of actions available to them whenever the first contracted period is reached.[12] Formally, a contract can be specified as $C = (C_1, C_2)$ with $C_1 = C_2$ where $\tilde{A}_i(h) = C_i(h) \in (\mathcal{P}(A_i) \setminus \emptyset)$ is the set of actions available to player i at history h. In addition, we impose that contracts are symmetric.[13]

Contracting comes at lump sum *enforcement costs* $\kappa \geq 0$ that players incur once at the first contracted period defined as $t_\kappa := \min\{t : C_t \neq A_t\}$ (rather than at the contracting stage).[14] These assumptions make our model more tractable and eliminate potential biases in favor of shorter and simpler smooth-landing contracts.[15]

We illustrate the timing of the contracting in Figure 1. In this section, we establish the conditions for full cooperation with contractual commitment for three cases: (a) full commitment, in which players commit to the cooperative action in every period (i.e., the contract is enforced in $t = 0$ at cost κ, with a present-discounted value of costs κ, and

[11] Later in the paper, we relax this pre-commitment assumption and allow for endogenous contract timing.

[12] This is the same as saying that the contract entered in $t = 0$ does not restrict the action space nor does it prescribe any actions between the contract stage and the first contracted period.

[13] Symmetry is in the sense that $C_i(h_i, h_j) = C_j(h_j, h_i)$ where (h_i, h_j) is a decomposition of history h into every action taken by any one player.

[14] Enforcement costs are associated with the costs of enforcing a contract (e.g., Schwartz and Watson, 2004) or the costs of monitoring compliance and detecting violations (e.g., Diamond, 1984, 1991). In a later extension, we will consider more general contracting costs that also include drafting costs (incurred at the time the contract is entered) associated with the costs that arise from researching and analyzing contingencies (*search costs*) (e.g., Klein, 2002; Tirole, 2009) and from specifying these contingencies (*ink costs*) (e.g., Dye, 1985; Anderlini and Felli, 1994; Melumad et al., 1997).

[15] The lump sum costs are independent of the length of the contract, the number of contracted periods, the amount of committed actions, or any other contingency.

Figure 1: Timeline of Contracting

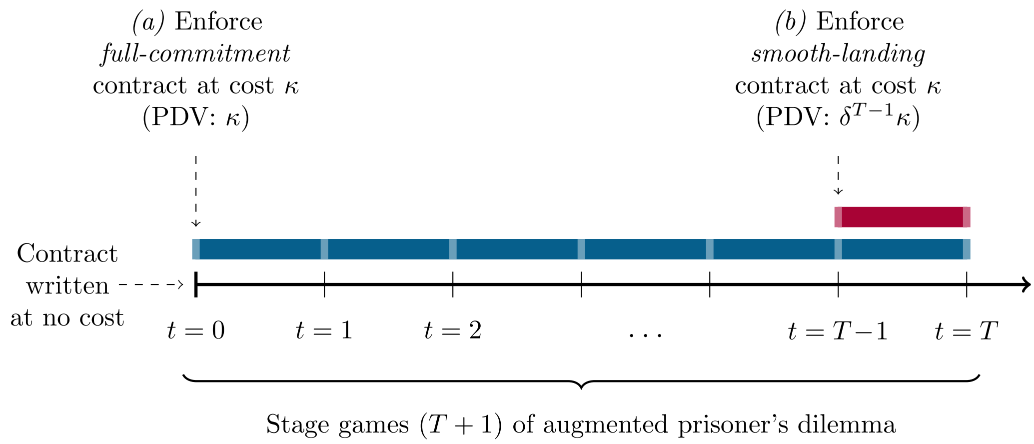

governs the business relationship in all t), and *abridged* full commitment as a special case; (b) a *simple* smooth-landing contract, in which committed actions are not contingent on the history of actions prior to the contract; and (c) an *improved* smooth-landing contract, in which the contracted actions depend on the history leading up to the contract. Both types of smooth-landing contracts are enforced in a period *after* the time of drafting and govern the business relationship in all subsequent periods. For instance, in Figure 1, we depict the case of contract enforcement in $T - 1$, at a present-discounted value of costs $\delta^{T-1}\kappa$.

3.1 Full-Commitment Contracts

The case of full commitment is the polar opposite of the environment without any commitment in our baseline scenario. Both players commit to a restricted action space $\tilde{A}_i$ for every period. Trivially, a full-commitment contract that restricts the action space to $\tilde{A}_i = \{w\}$, thus effectively prescribing $a_{i,t} = w$ for all i and t, implements full cooperation for the entirety of the extended game and for any discount factor δ.

A potential shortcoming of the full-commitment contract, particularly when compared to the partial-cooperative baseline equilibrium, is that it comes into effect during periods when the baseline already achieves cooperation. Because costs are paid as enforcement costs at the time the contract becomes active, a contract that is active for the entire duration of the game is not cost-effective.

A shorter alternative confines the action space to $\tilde{A}_i = \{w\}$, but, unlike its longer counterpart, is only active for the last τ periods of the game (for instance, $\tau = 2$, as the smooth-landing contract depicted in Figure 1). Such an *abridged* full-commitment contract is not effective in preventing the unraveling present in the baseline; it rather shifts it forward in time. This is because such a contract partitions the game into an augmented prisoner's dilemma with $T + 1 - \tau$ periods played first, followed by τ periods of contracted cooperation (i.e., abridged full-commitment contract for the last periods). The end-of-game problem thus moves from the last stage game to the τ-last stage, and the players choose (l, l) in period $T+1-\tau$ in equilibrium (i.e, the period before the τ-last stage).

Proposition 1 (Abridged Full-Commitment Contract)**.** *An abridged full-commitment contract is unable to sustain the cooperative outcome for the entirety of the relationship.*

Note that the cooperation threshold $\bar{\delta}^{BL}$ in Lemma 1 is independent of the length of the game. Hence, the partitioned augmented prisoner's dilemma will feature a partial-cooperative equilibrium in addition to the non-cooperative one, as laid out in Lemmas 1 and 2.

3.2 Hybrid Contracting

Saying that the abridged full-commitment contract does not work in achieving the desired cooperative outcome in every period does not mean that shorter contracts are ineffective in general. Rather, by being entirely rigid when in effect, the abridged full-commitment contract eliminates the possibility of rewarding or punishing players based on their previous behavior. To rectify this deficiency, we introduce a class of contracts that we refer to as *smooth-landing contracts*. They are in effect for less than the entire duration of the game and maintain the ability to react to previous actions and potential deviations from the cooperative behavior. This flexibility of smooth-landing contracts results in a *hybrid contract*, in which parties formally commit to transactional boundaries while retaining the relational-contract dynamics that have governed the early stages of their relationships—similar to formal relational contracts discussed in Frydlinger et al. (2019). An abridged (and rigid) full-commitment contract, on the other hand, no longer resembles a relational contract.

3.2.1 Simple Smooth-Landing Contracts

In a simple version of the smooth-landing contract, the players commit to an action space that effectively prohibits some actions for the latter part of their interaction.[16] In the simple case, the action space $\tilde{A}_i$ does not depend on past actions.[17] For now, we consider contracts that take effect only in the last period T (so that $\tau = 1$), without interfering with the players in the preceding periods ($t = 0, \ldots, T-1$). Such a "smooth-landing contract" then prohibits the lowest-effort action s by specifying a (restricted) action space $\tilde{A}_i = \{l, w\}$.

[16] We assume players abide by the rules of the contract, which means they do not choose the eliminated action. This can be achieved by including appropriate penalties. We do not specify such contracted penalties but simply assume that a violation of the contract terms is prohibitively costly.

[17] As discussed before, we require the restrictions on the action space to be symmetric.

Table 2: Restricted Stage-Game Under Simple Smooth-Landing Contract

	loaf (l)	work (w)
loaf (l)	(π^L, π^L)	(π_{wl}, π_{lw})
work (w)	(π_{lw}, π_{wl})	$(1, 1)$

Definition 1. *The simple smooth-landing contract specifies for $t = T$: $\tilde{A}_i(h) = \{l, w\}$ for all h.*

With this contract, the (restricted) stage game (in the last period of the repeated game) is summarized in Table 2. By Condition 1, we have $\pi_{wl} < 1$ and $\pi_{lw} < \pi^L$, and the restricted game continues to have two pure-strategy Nash equilibria, (l, l) and (w, w). Unlike in the three-action augmented prisoner's dilemma in Table 1, however, the cooperative outcome (w, w) is now one of these equilibria.

Simple Smooth-Landing Strategy. *If all parties have cooperated without failure, continue to cooperate and play w. If a party has deviated at any point, play s (as a punishment) until the penultimate period $T - 1$. In the last period, punish deviations by playing l. Formally, this is saying every party plays the modified grim trigger strategy*

$$a_{i,t}^{SL} = \begin{cases} s & \textit{if } t \neq T \wedge \exists t' < t : a_{j,t'} \neq w; \\ l & \textit{if } t = T \wedge \exists t' < t : a_{j,t'} \neq w; \\ w & \textit{otherwise.} \end{cases} \tag{5}$$

This strategy is analogous to the baseline strategy, except for the following point: In the last period, the contract rewards earlier cooperation with (w, w) and punishes a deviation with (l, l) (both outcomes are Nash equilibria in the restricted stage-game in Table 2). Under this strategy, cooperation is a player's preferred choice in $t = T$ and all earlier periods whenever

$$\delta \geq \frac{\pi^D - 1}{1 - \pi^L} =: \bar{\delta}^{SL}. \tag{6}$$

We summarize this result of cooperation under the simple smooth-landing contract in the following lemma:

Lemma 3. *For a simple smooth-landing contract, a fully cooperative equilibrium, yielding the outcome $(w, w), \ldots, (w, w)$, exists whenever $\delta \geq \frac{\pi^D - 1}{1 - \pi^L} =: \bar{\delta}^{SL}$.*

Through a simple restriction of the action space, our smooth-landing contract preserves the ability to reward cooperative behavior and to punish non-cooperative behavior in the penultimate stage $T - 1$. This ability is crucial and is missing from the abridged full-commitment contract in Proposition 2. Put differently, smooth-landing contracts that are not themselves contingent only work if not too rigid (a result paralleling that in Battigalli and Maggi (2002)) by retaining some of the relational dynamics that have governed the parties' relationship in earlier periods.

3.2.2 Improved Smooth-Landing Contracts

The simple smooth-landing contract is simple in the sense that it is not contingent and leaves punishment to the players' strategies. As a consequence, punishment for deviations as well as reward for cooperation in the last period must form a Nash equilibrium in the stage game. In our baseline case, that is action l as a reward and s as punishment in the grim trigger strategy.[18] The simple smooth-landing contract modifies the last period stage game such that (w, w), the payoff-maximizing cooperative outcome, forms a Nash equilibrium of the restricted stage game, at the cost of having to punish deviation with the less harsh (l, l) equilibrium.

However, when using a contingent contract, rewards that are not part of a Nash equilibrium strategy can be leveraged, which in turn allows the players to render deviations relatively more costly. In our case, the players can achieve this by contracting action w also at the last stage (provided there are no prior deviations) while continuing to punish with s in the case of prior deviations. This contract, therefore, rewards uninterrupted cooperation and punishes (in the last period) any deviations.[19] The corresponding contract is defined below:

Definition 2. *An improved smooth-landing contract specifies for $t = T$: $C_i(h) = \{w\}$ if $h = (w, w)^{T-1}$ and $C_i(h) = \{s\}$ otherwise.*

This improved contract type then requires the usage of an adapted grim trigger strategy described below.

Improved Smooth-Landing Strategy. *If all parties have cooperated without failure, continue to cooperate and play w. If a party has deviated at any point, play s (as a punishment) for the remainder of the game. Formally, this is saying every party plays the modified grim trigger strategy*

$$a_{i,t}^{SL^+} = \begin{cases} s & \text{if } \exists t' < t : a_{j,t'} \neq w; \\ w & \text{otherwise.} \end{cases} \tag{7}$$

We achieve full cooperation in the last period by prescribing the cooperative action in the last period through the improved smooth-landing strategy. Taking this outcome

[18] Note from equation (2) that, except for the last period, the standard grim trigger strategy $a_{i,t}^{GT}$ already implements the most effective (i.e., the payoff-maximizing) reward.

[19] Note that this is not saying that this contingent contract chooses the harshest punishment. We discuss this in more detail when comparing the contracts in terms of optimality in Section 4.4.

of full cooperation in $t = T$, a deviation in $T-1$ is then unprofitable and cooperation is sustainable whenever

$$\delta \geq \pi^D - 1 =: \bar{\delta}^{SL^+}. \tag{8}$$

We summarize this result of cooperation under the improved smooth-landing contract in the following lemma:

Lemma 4. *For an improved smooth-landing contract, a fully cooperative equilibrium in which $a_{i,t}^{SL^+}$ is played, yielding the outcome $(w,w),\ldots,(w,w)$, exists whenever $\delta \geq \pi^D - 1 =: \bar{\delta}^{SL^+}$.*

Unlike in the simple smooth-landing contract, where the cooperative action is merely one of the allowed actions in the parties' choice set, the improved smooth-landing contract incorporates the cooperative action w as the contracted last-period action. However, unlike in the abridged full-commitment contract, the prescription of the cooperative action in the last period is contingent on the parties' actions in earlier periods. The improved smooth-landing contract thus retains the ability to punish deviations in earlier periods. With a more costly punishment than in the simple case, the improved smooth-landing contract can implement the full-cooperative outcome just like the simple contract, but for a broader range of discount factors δ because $\bar{\delta}^{SL^+} < \bar{\delta}^{SL}$.

4 Comparative Analysis and Contract Choice

In this section, we turn to a comparative analysis of the different hybrid-contracting scenarios. *First*, we start with *implementability* of the cooperative outcome by summarizing the conditions under which the fully-cooperative outcome (with commitment) and the partial-cooperative outcome (without commitment) can be achieved. *Second*, given implementability, we compare the payoffs under the scenarios with commitment (full-commitment contract and smooth-landing contracts) with the payoffs in the baseline to establish the *profitability* of hybrid contracting relative to the baseline (that relies solely on relational dynamics). *Third*, we determine the parties' choice of contract among the profitable alternatives. We then conclude this section with a discussion of optimal hybrid contracting.

Smooth-landing contracts affect cooperation at two different margins. First, at the *intensive margin*, the use of an effective smooth-landing contract extends the number of periods in which the agents can reach the cooperative outcome, from T in the partial-cooperative baseline equilibrium (Lemma 1) to $T+1$ in the equilibrium with a smooth-landing contract (Lemmas 3 and 4). Second, at the *extensive margin*, a smooth-landing contract is more forgiving and tolerates lower levels of patience (measured by δ) to reach

Table 3: Summary of Implementability Thresholds for Contract Classes

	Implementability threshold	**Existence condition**
Baseline (cooperative eqm.)	$\bar{\delta}^{BL} = \frac{\pi^D - 1}{\pi^L} = \frac{4\pi - 2m - 2}{\tilde{e}}$	$\tilde{e} \leq 4\pi - 2m - 2$
Simple smooth-landing contr.	$\bar{\delta}^{SL} = \frac{\pi^D - 1}{1 - \pi^L} = \frac{4\pi - 2m - 2}{2 - \tilde{e}}$	$\tilde{e} \leq 4 + 2m - 4\pi$
Improved smooth-landing contr.	$\bar{\delta}^{SL^+} = \frac{\pi^D - 1}{1 - 0} = 4\pi - 2m - 2$	$m \geq 2\pi$

Notes: This table summarizes the implementability thresholds $\bar{\delta}^k$ stated in Lemma 1 (partial-cooperative baseline equilibrium), Lemma 3 (simple smooth-landing contract), and Lemma 4 (improved smooth-landing contract), as well as the existence conditions (so that $\bar{\delta}^k \leq 1$).

the cooperative outcome in each period. This means that a contract can also induce cooperation by changing the equilibrium from a non-cooperative one in the baseline (Lemma 2) to a cooperative equilibrium under a smooth-landing contract.

4.1 Implementability of Cooperative Equilibria

We summarize the respective implementability thresholds for the partial-cooperative equilibrium in the baseline ($\delta \geq \bar{\delta}^{BL}$), the simple smooth-landing contract ($\delta \geq \bar{\delta}^{SL}$), and the improved smooth-landing contract ($\delta \geq \bar{\delta}^{SL^+}$). The existence condition ensures that $\bar{\delta}^k \leq 1$ (where k denotes the different scenarios). We depict these conditions and the range of parameters for which they are satisfied in Figure 2 for a specific parameterization of our model.

We make three simple observations. First, all thresholds depend positively on π^D, meaning that increasing the payoff of deviating makes any form of cooperation harder to sustain. Moreover, because $\pi^D = 2\pi - m$, an increase in m reduces the thresholds, because higher moral costs of free-riding reduce the attractiveness of deviating from the cooperative equilibrium.

Second, the role of $\pi^L = \tilde{e}/2$, the payoff of the medium-effort outcome (i.e., the loafing outcome), is more nuanced. The loafing equilibrium plays a different role in each of our scenarios, acting as the final reward in the baseline and as the final punishment in the simple smooth-landing contract. In the improved smooth-landing contract, however, it plays no role. As a direct consequence, increasing π^L has no effect on the improved smooth-landing contract, whereas the simple smooth-landing contract's threshold increases, because the punishment becomes less severe. In contrast, in the partial-cooperative baseline, an increase in π^L results in an increase of the cooperation reward, lowering the respective threshold.

Third, because the intermediate payoff π^L is directly proportional to the intermediate effort level $\tilde{e}$, the above dynamic extends to changes in $\tilde{e}$. Moreover, this implies that

Figure 2: Existence of Cooperative Equilibria

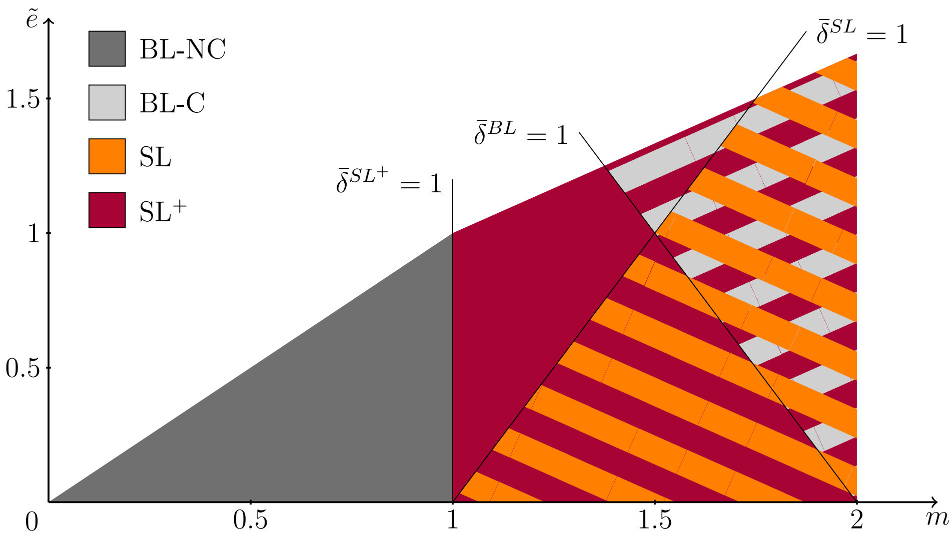

Notes: The figure depicts (in m-$\tilde{e}$ space for $T = 3$ and $\pi = 1.5$) the regions in which full cooperation in all t can be achieved with simple (SL, orange) and improved (SL$^+$, red) smooth-landing contracts. We also depict regions in which the partial-cooperative baseline equilibrium can be implemented (BL-C, lightgray)..

m and $\tilde{e}$ act as substitutes in the baseline and as complements for the simple smooth-landing contract, resulting in the decreasing existence frontier for $\bar{\delta}^{BL}$ and the increasing existence frontier for $\bar{\delta}^{SL}$ as depicted in Figure 2.

While all smooth-landing contracts require a sufficiently high discount factor for the fully cooperative equilibrium, the full-commitment contract is a special case. Because players commit to one action for the entire game with no opportunity to deviate, full-commitment contracts are, by design, able to implement the fully cooperative equilibrium for any δ and any $(m, \tilde{e})$-parametrization in Figure 2 (and are thus omitted).

4.2 Profitability of Contracts

Given that a fully cooperative equilibrium is implementable, a contract is said to be profitable if the net payoff's present-discounted value is higher than under the partial-cooperative or non-cooperative baseline equilibrium. Recall that the partial-cooperative equilibrium of the baseline model is characterized by a payoff stream of $\pi_i(w, w) = 1$ in the first T periods and $\pi_i(l, l) = \pi^L$ in the $(T + 1)^{\text{th}}$. The non-cooperative equilibrium, by contrast, pays out $\pi_i(l, l) = \pi^L$ in every period. Recall further that we model contractual costs as enforcement costs κ paid in the first period in which the contract comes into effect.

Full-Commitment Contract. The full-commitment contract yields a stage-payoff of $\pi_i(w, w) = 1$ for $T + 1$ periods and is *ex-ante* preferable to the partial-cooperative equilibrium (Lemma 1) whenever

$$\frac{1-\delta^T}{1-\delta} + \delta^T - \kappa \geq \frac{1-\delta^T}{1-\delta} + \delta^T \pi^L \iff \kappa \leq \delta^T \left(1-\pi^L\right). \tag{9}$$

Intuitively, because the full-commitment contract changes only the payoff of the last period, the cost of the contract can only be as high as the discounted difference in the last-period payoffs.

For the full-commitment contract to payoff-dominate the non-cooperative equilibrium (Lemma 2), we need

$$\frac{1-\delta^T}{1-\delta} + \delta^T - \kappa \geq \left(\frac{1-\delta^T}{1-\delta} + \delta^T\right) \pi^L \iff \kappa \leq \left(\frac{1-\delta^{T+1}}{1-\delta}\right) \left(1-\pi^L\right). \tag{10}$$

This cost threshold is higher than for the partial-cooperative equilibrium. This also means that there are contracting costs κ for which full-commitment contracts are only preferred if the alternative is a non-cooperative equilibrium. We summarize these results in the following proposition:

Proposition 2. *For $\kappa \leq \frac{1-\delta^{T+1}}{1-\delta}\left(1-\pi^L\right)$, the full-commitment contract dominates the non-cooperative baseline. For $\kappa \leq \delta^T \left(1-\pi^L\right)$, the full-commitment contract dominates both the partial-cooperative and non-cooperative baseline equilibria.*

(Simple and Improved) Smooth-Landing Contract. Both types of smooth-landing contracts (when implementable, that is, when implementing the fully cooperative equilibrium) yield a payoff stream of $\pi_i(w, w) = 1$ in every period. The contracts, however, come into effect only in the final period of the game, and enforcement costs are incurred only in that period. As a consequence, smooth-landing contracts are profitable compared to the partial-cooperative baseline (Lemma 1) when

$$\frac{1-\delta^T}{1-\delta} + \delta^T (1-\kappa) \geq \frac{1-\delta^T}{1-\delta} + \delta^T \pi^L \iff \kappa \leq 1-\pi^L, \tag{11}$$

or when the costs (incurred in the final period) are not higher than the payoff difference (materializing in that period). Moreover, a smooth-landing contract is profitable relative to the non-cooperative baseline equilibrium (Lemma 2) when

$$\frac{1-\delta^T}{1-\delta} + \delta^T (1-\kappa) \geq \frac{1-\delta^{T+1}}{1-\delta} \pi^L \iff \kappa \leq \frac{1-\delta^{T+1}}{(1-\delta)\,\delta^T} \left(1-\pi^L\right). \tag{12}$$

This cost threshold is higher than for the partial-cooperative equilibrium case. We summarize the results for both types of smooth-landing contracts in the following proposition:

Proposition 3. *For* $\kappa \leq \frac{1-\delta^{T+1}}{(1-\delta)\delta^T}\left(1-\pi^L\right)$*, an implementable smooth-landing contract dominates the non-cooperative baseline equilibrium. For* $\kappa \leq 1-\pi^L$*, it dominates both the partial-cooperative and non-cooperative baseline equilibria.*

4.3 Contract Choice

To determine which contract is optimal from the players' *ex-ante* perspective, we compare the net payoffs of the different contract types: the contract featuring full commitment, the simple smooth-landing contract, and the improved smooth-landing contract. A comparison with the baseline also shows when a no-commitment scenario dominates any of the contractual solutions and the ensuing hybrid-contracting strategy. We refer to an implementable smooth-landing contract as one that implements the fully cooperative equilibrium (Lemmas 3 and 4). We summarize our contract-choice results in Proposition 4.

Proposition 4. *A comparison between the two types of smooth-landing contracts, the full-commitment contract, and the baseline admits five parameter regions:*

1. *for* $\delta \in [0, \bar{\delta}^{SL^+}]$*, the full-commitment contract dominates both smooth-landing contract types and the non-cooperative baseline for sufficiently low* κ *(otherwise agents play the non-cooperative baseline);*
2. *for* $\delta \in [\bar{\delta}^{SL^+}, \min\left\{\bar{\delta}^{SL}, \bar{\delta}^{BL}\right\}]$*, the improved smooth-landing contract is implementable and dominates the full-commitment contract; given a sufficiently low* κ*, it also dominates the baseline (otherwise agents play the non-cooperative baseline);*
3. *for* $\delta \in [\bar{\delta}^{SL}, \bar{\delta}^{BL}]$*, if* $\bar{\delta}^{SL} < \bar{\delta}^{BL}$*, both smooth-landing contract types are implementable, identical in payoffs, and dominate the full-commitment contract; given a sufficiently low* κ*, they also dominate the baseline (otherwise agents play the non-cooperative baseline);*
4. *for* $\delta \in [\bar{\delta}^{BL}, \bar{\delta}^{SL}]$*, if* $\bar{\delta}^{SL} > \bar{\delta}^{BL}$*, the improved smooth-landing contract is implementable and dominates the full-commitment contract; given a sufficiently low* κ*, it also dominates the baseline (otherwise agents play the partial-cooperative baseline);*
5. *for* $\delta \in [\max(\bar{\delta}^{SL}, \bar{\delta}^{BL}), 1]$*, both smooth-landing contract types are implementable, identical in payoffs, and dominate the full-commitment contract; given a sufficiently low* κ*, they also dominate the baseline (otherwise agents play the partial-cooperative baseline).*

In Figure 3, we plot a player's profits $\pi_i(\cdot)$ against the discount factor δ, and thus provide an illustration of the results in Proposition 4. For low values of κ in panel (a), we observe that the full-commitment contract dominates the non-cooperative baseline

Figure 3: Payoff Comparison between Equilibria

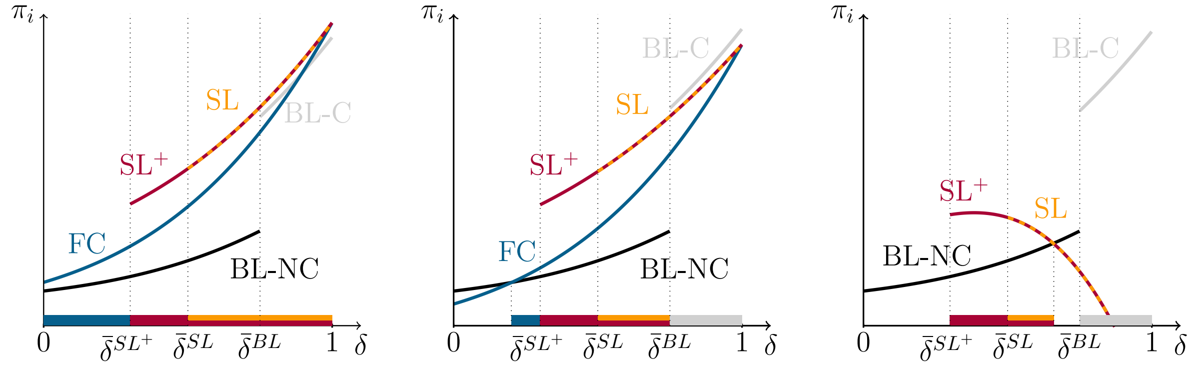

(a) $\kappa = 0.5$ (b) $\kappa = 0.7$ (c) $\kappa = 5$

Notes: The figures depict a player i's payoffs $\pi_i(\cdot)$ under the different contracting choices from the different contracts (for $T = 3$, $\pi = 1.5$, $m = 1.7$, $\tilde{e} = 0.8$, and selected values of κ). BL-C and BL-NC denote the partial-cooperative and non-cooperative equilibrium payoffs of the baseline, respectively. The colored bars at the bottom indicate which contract is preferred for the specific δ-region (blue for the full-commitment contract, orange for the simple smooth-landing contract, and red for the improved smooth-landing contract).

outcome in all of parameter region 1 of the proposition, where a smooth-landing contract is not implementable (see the blue-shaded bar). For higher values of κ (in panels (b) and (c)), the dominance of the full-commitment contract is limited. In parameter region 2, only the improved smooth-landing contract is implementable, and its payoffs are higher than for the full-commitment contract and the non-cooperative baseline (red-shaded bar). In parameter regions 3 through 5, both types of smooth-landing contracts are implementable and generate the same profits (because, for a given δ, they yield the same outcome at the same contracting cost). They always dominate the full-commitment contract, but players prefer them to the baseline only for sufficiently small contracting costs κ (red/orange-shaded bars).

Figure 4 provides an alternative visualization of the results in Proposition 4, depicting the contract choice for varying values of the discount factor δ and contracting costs κ (and our example parameterization). First, note that for $\delta > \bar{\delta}^{BL} = 0.75$, a partial-cooperative equilibrium is sustainable in the baseline. If, additionally, $\kappa < 1 - \pi^L = 0.6$, both smooth-landing contract types yield a higher payoff than the baseline. In this case, a smooth-landing contract improves cooperation at the *intensive margin*. In contrast, for $\delta < \bar{\delta}^{BL}$, the baseline features only a non-cooperative equilibrium. Consequently, any contract that induces cooperation does so at the *extensive margin*. For our example parameterization of the model, we have $\bar{\delta}^{SL} = 0.5 > 0.3 = \bar{\delta}^{SL^+}$, and there exists a low range of discount factors for which only an improved smooth-landing contract can implement full cooperation. Players incur costs κ in the last period and discount these costs accordingly. As a consequence, a lower δ decreases the effective costs of the contract,

Figure 4: The Choice of (Hybrid) Contracting

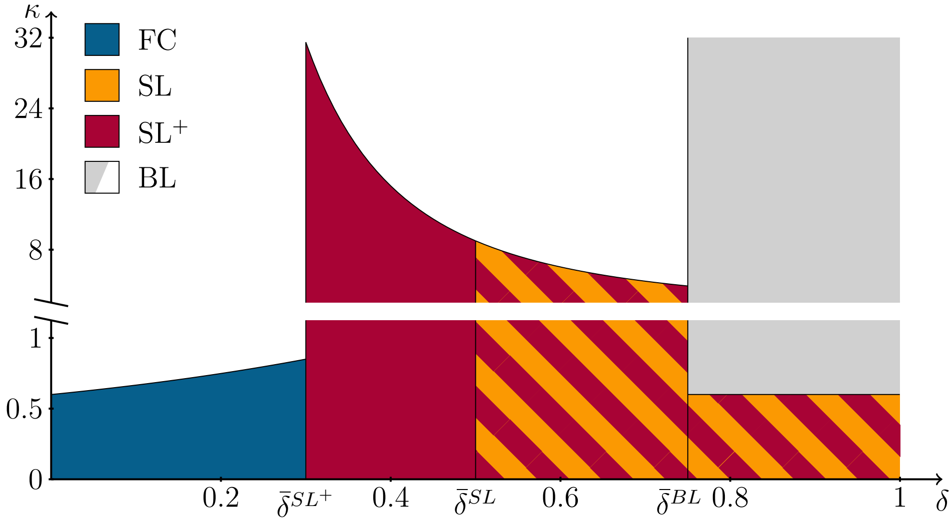

Notes: The figure depicts (in δ-κ space for $T = 3$, $\pi = 1.5$, $m = 1.7$, and $\tilde{e} = 0.8$) the respective regions with the highest-payoff contracts. In red-shaded regions, the improved smooth-landing contract is the preferred (and only) choice; in red/orange-shaded regions, players are indifferent between the two types of smooth-landing contracts; in the blue region, the full-commitment contract is chosen; in the light-gray region, the partial-cooperative baseline outcome dominates all contracts.

implying an increase in the cost threshold below which a smooth-landing contract still dominates—explaining the decreasing cost frontier for intermediate values of δ. On the other hand, the cost for the full-commitment contract is always paid in the first period and does not depend on variations in the discount factor. However, while costs remain constant, the relative utility of cooperation induced in the last period decreases with a lower discount factor, resulting in a decrease in the cost threshold below which full commitment dominates the non-cooperative baseline outcome—explaining the increasing cost frontier for low values of δ.

4.4 Optimal Hybrid Contracting

Within our class of smooth-landing contracts, we have assumed two straightforward types. Whenever these are *implementable*, they implement the optimal outcome; and whenever they are profitable, they are the lowest-cost contractual solution. A one-period smooth-landing contract is, therefore, the best (or *optimal*) hybrid contracting solution, implementing the full-cooperation outcome at the lowest possible cost.[20] No other contract or contract class with enforcement costs κ incurred in T can do better.

[20]Note that, by assumption of positive contracting costs $\kappa > 0$, an optimal contract is optimal only in a second-best sense (whereas the first-best outcome is cooperation in all periods without a contract and thus at no cost). Because of the end-of-game properties of our finitely-repeated stage game, the equilibrium in a commitment-free scenario cannot be first-best.

Proposition 5. *Both types of smooth-landing contracts are optimal whenever they are implementable and profitable.*

When a smooth-landing contract is not implementable, a better contractual arrangement may exist; one that is implementable for a wider range of values of δ. For instance, recall that the improved smooth-landing contract does not improve upon the simple smooth-landing contract in terms of its payoffs. However, by rendering the contract dependent on past actions, the improved smooth-landing contract is able to punish deviations (off-equilibrium) more effectively. The implementability threshold of every contract depends on the gap between the reward for cooperation and the punishment for deviation. Because under full cooperation, the reward for cooperation is the same for all contracts, one can find the optimal (one-period) smooth-landing contract (with the lowest possible implementability threshold) by optimizing the (off-equilibrium) punishment.

Beyond one-period contracts, multi-period contracts can prolong punishment and thus render the contracts more effective. Due to earlier enforcement costs, longer contracts incur higher effective costs, but they can be optimal because they enable full cooperation when a one-period contract fails. We discuss such multi-period contracts as our first extension in the next section.

5 Extensions

In this section, we relax three assumptions from our main analysis. First, we allow for multi-period smooth-landing contracts. By allowing for harsher punishment of deviations, this first scenario extends the capability of smooth-landing contracts to implement full cooperation (i.e., more cooperation at the extensive margin) to previously excluded regions. For the second and third extensions, we consider more general contracting cost structures and allow players to enter the smooth-landing contract at a later stage (rather than the outset of their relationship). Both extensions show that the results in the main model are robust to these generalizations.

5.1 Multi-Period Smooth-Landing Contracts

In the baseline model, we assume a single-period smooth-landing contract with $\tau = 1$. We show that for sufficiently high discount factors δ, this single-period contract can implement the full-cooperation outcome by ensuring cooperation in the very last period (where the partial-cooperative baseline equilibrium fails). An effective multi-period smooth-landing contract, therefore, will not generate more cooperation at the *intensive margin* because the partial-cooperative baseline equilibrium yields cooperation in all but the last period (and therefore in some periods the multi-period smooth-landing contract covers). Moreover, for a multi-period contract, the parties will incur the contracting costs earlier, implying higher discounted costs. We show below that, while not yielding any improvements at the intensive margin, multi-period smooth-landing contracts can generate more cooperation (relative to single-period contracts) at the *extensive margin* by lowering the threshold $\bar{\delta}^k$.

Consider a smooth-landing contract as described in Section 3, where the prescription is in effect for a total of $\tau \geq 2$ periods—for instance, the smooth-landing contract depicted in Figure 1 is with $\tau = 2$. The contract ensures cooperation in these last τ periods (with contract enforcement in $T + 1 - \tau$), and it implements the cooperative outcome if the player cooperates in the last period, $T - \tau$, before contract enforcement (and all earlier periods). The larger τ and the earlier a player takes this decision to deviate (or not), the longer is the stream of punishment payoffs (π^L or 0) and the higher are the opportunity costs from a deviation, that is, while the immediate benefits from a deviation are constant. Deviation becomes less likely as the player will now—that is, for longer contracts—cooperate for even lower values of δ. The reasoning is analogous to that for the partial-cooperative baseline equilibrium in Lemma 1 where the strongest deviation incentives are in the penultimate period.

The following lemma summarizes the relationship between the implementability threshold $\bar{\delta}^k$ and the contract duration τ.

Lemma 5. *The implementability thresholds $\delta^k(\tau)$ for the simple ($k = SL$) and improved ($k = SL^+$) smooth-landing contracts are decreasing in τ for all $\tau \leq \bar{\tau}^k$ with $\bar{\tau}^k$ the largest τ such that $\bar{\delta}^k(\tau) \geq \frac{\pi^D - 1}{\pi^D} = \bar{\delta}^D$. These implementability thresholds are constant for all larger values of τ.*

Longer smooth-landing contracts are more costly to implement full cooperation for a given δ, but longer contracts can lower the implementability threshold up to $\bar{\tau}^k$. For any higher τ, players with a discount factor below the implementability threshold are so impatient that they would always prefer to deviate at the first rather than the last possible point in time. For these agents, there is no smooth-landing contract that can provide an effective punishment. This means that a multi-period contract beyond this threshold triggers higher costs without further improvements at the extensive margin—imposing a natural upper bound on the length of the smooth-landing contract.

Proposition 6. *The length τ of an implementable multi-period smooth-landing is bounded above by* $\min\left\{T, \bar{\tau}^k\right\}$.

For a given δ, multi-period smooth-landing contracts implement the same payoffs as single-period contracts for a higher cost, but they can extend the circumstances in which full cooperation can be achieved. The choice of contract has to balance this trade-off between the intensive margin and the extensive margin. A multi-period contract is chosen only when a shorter contract is no longer implementable (when δ is too low).

Figure 5: Contract Choice under Multi-Period Contracting

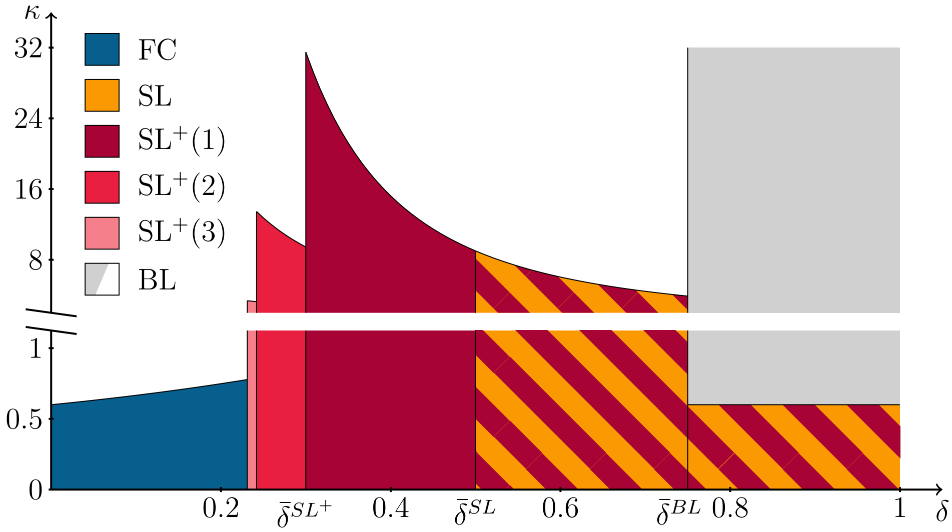

Notes: The figure depicts (in δ-κ space for $T = 3$, $\pi = 1.5$, $m = 1.7$, and $\tilde{e} = 0.8$) the respective regions with the highest-payoff contracts. In red-shaded regions, the improved smooth-landing contract is the preferred (and only) choice; in red/orange shaded regions, players are indifferent between the two types of smooth-landing contracts; in the blue region, the full-commitment contract is chosen; in the light-gray region, the partial-cooperative baseline outcome dominates all contracts. The lighter shades of red capture results for smooth-landing contracts with a runtime of two ($\tau = 2$) and three periods ($\tau = 3$), respectively.

Moreover, an improved smooth-landing contract can be shorter than a simple smooth-landing contract to implement full cooperation for the same values of δ. In other words, for some δ, a simple smooth-landing contract is implementable only for $\tau = 2$, whereas an improved smooth-landing contract may still be implementable for $\tau = 1$, in which case the shorter improved smooth-landing contract is always chosen.

We depict the players' contract choice in Figure 5. As in Figure 4, the dark-red shaded region depicts the parameters in which a one-period (improved) smooth-landing contract is chosen. In the lighter-red-shaded regions, two-period and three-period (improved) smooth-landing contracts are chosen. The players choose a two-period contract for values of δ below the one-period threshold. The two-period contract can implement the full-cooperation outcome, but its choice is optimal only for lower values of κ. Similarly, for a three-period contract.

5.2 Contracting with Upfront Costs

The contracting costs in our main analysis are enforcement costs (κ_i), which players incur when a contract is enforced but not when it is drafted. In this extension, we allow for parts of the contracting costs κ to be upfront costs (κ_d) that arise when the parties draft

Table 4: Cost Thresholds for Implementable Smooth-Landing Contracts

	Partial-cooperative baseline	**Non-cooperative baseline**
Smooth-landing contract	$\kappa \leq \frac{\delta^T}{1-\alpha+\alpha^T}\left(1-\pi^L\right)$	$\kappa \leq \frac{1-\delta^{T+1}}{(1-\delta)(1-\alpha+\alpha^T)}\left(1-\pi^L\right)$
Full commitment	$\kappa \leq \delta^T(1-\pi^L)$	$\kappa \leq \frac{1-\delta^{T+1}}{1-\delta}\left(1-\pi^L\right)$

Notes: This table provides the cost thresholds (maximum total contracting costs κ) for the profitability of smooth-landing contracts. Total contracting costs are $\kappa = \kappa_d + \kappa_i$ with $\kappa_d = (1-\alpha)\,\kappa$ and $\kappa_i = \alpha\kappa$.

the contract (in $t = 0$). Overall costs are then

$$\kappa = \kappa_d + \kappa_i \tag{13}$$

with drafting costs $\kappa_d = (1-\alpha)\,\kappa$ and enforcement costs $\kappa_i = \alpha\kappa$. Our analysis in the previous sections is for the scenario with $\alpha = 1$.

The contracting-cost structure does not affect the ability of smooth-landing contracts to implement full cooperation (in Lemmas 3 and 4)—because given a contract is in place and enforced, the contract costs are sunk. Profitability, however, is a function of contracting costs. We reformulate the profitability thresholds for the simple and improved smooth-landing contracts and summarize them in Table 4. Note that the profitability of a full-commitment contract is not driven by the contracting-cost structure (nor are the thresholds), because both drafting and enforcement costs arise at $t = 0$.

In Figure 6, we depict the optimal contract choice under three cost regimes (where enforcement costs make up (a) none, (b) half, and (c) three-quarters of total costs). Comparing these contract choices with those in the main analysis (in Figure 4) reveals that a more flexible cost structure does not change the qualitative nature of our findings—although the exact regions in which different contract types are chosen are different.

First, recall that upfront contracting costs do not affect the profitability of a full-commitment contract (blue-shaded region in Figure 6) because parties incur both drafting and enforcement costs at the outset of their relationship. When $\alpha = 0$, for smooth-landing contracts, too, all contracting costs are due at the outset of the game (rather than at the enforcement stage later). Then, whenever a smooth-landing contract is implementable, it is payoff-equivalent to a full-commitment contract; the shaded regions in panel (a) overlap. As α increases, the picture changes. For smooth-landing contracts, a smaller fraction of the costs is paid upfront, and a larger fraction later. These contracts, therefore, become effectively cheaper. Whenever a smooth-landing contract is implementable (and the partial-cooperative baseline equilibrium is not), higher values of α mean that the contract is optimal for higher total contracting costs κ. We can see this in panels (b) and (c) of Figure 6.

Figure 6: Smooth-Landing Contracts with Partial Upfront Costs

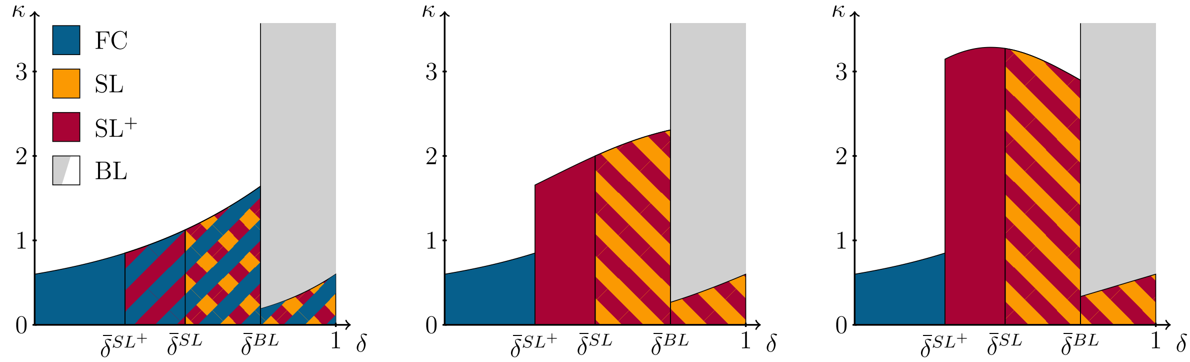

(a) Costs with $\alpha = 0$ (b) Costs with $\alpha = 0.5$ (c) Costs with $\alpha = 0.75$

Notes: The figure depicts (in δ-κ space for single-period smooth-landing contracts with $T = 3$, $\pi = 1.5$, $m = 1.7$, and $\tilde{e} = 0.8$) the respective regions with the highest-payoff contracts. In red-shaded regions, the improved smooth-landing contract is the preferred (and only) choice; in red/orange-shaded regions, players are indifferent between the two types of smooth-landing contracts; in the blue region, the full-commitment contract is chosen; in the light-gray region, the partial-cooperative baseline outcome dominates all contracts.

The cost frontier for smooth-landing contracts is decreasing when $\alpha = 1$ (in Figure 4) and increasing when $\alpha = 0$ (in panel (a) of Figure 6). For intermediate values, the frontier takes on a hump-shaped form. The reason is that both effects discussed in the context of Figure 4 factor in.[21]

The lower the value of α, the higher the fraction of the contracting costs paid upfront, the smaller the fraction paid as enforcement costs later in the relationship. Lower values of α thus increase the effective contracting costs. However, it is misleading to assume that contract parties should respond by negotiating the contract later to push the drafting costs further into the discounted future. As we demonstrate in the next extension, subsequent negotiations of the smooth-landing contracts undermine the parties' incentives to cooperate. While the parties save on contracting costs, their smooth-landing contract no longer implements the full-cooperation outcome.[22]

5.3 Endogenous Contract Timing

In our main model, the parties enter into contracts in the first period, prior to any interaction. We now extend the model such that in every period, players take two decisions: First, whether or not to enter a contract, and second, which action a_i to take. For the

[21]For high values of α, a lower δ decreases the effective costs of the contract and increases the cost threshold below which a smooth-landing contract dominates—implying a decreasing cost frontier. For low values of α, a larger fraction of the costs is paid up front. Lower values of δ reduce the relative utility of cooperation induced in the last period while only minimally reducing effective costs (as a larger fraction is not discounted)—implying an increasing cost frontier. When α takes intermediate values, the combined effects yield the hump-shaped frontier.

[22]The exposition in Section 5.3 is for $\alpha = 1$; the results, however, hold for all values of α.

analysis in this section, and for expositional ease, we restrict our attention to single-period improved smooth-landing contracts.

The crucial observation for solving the model under endogenous timing is that any agreed-upon contract can only be payoff-contingent on periods after the contract was negotiated. The reason for this is straightforward: If the contract were to discriminate against certain histories (that is, histories in which deviation has taken place), these histories would be followed by induced action profiles that yield lower outcomes than those following other histories. When the history in which the contract is drafted is reached, however, it is easy to see that it would be a profitable deviation to simply implement whichever actions are prescribed in the payoff-dominant history instead. Thus, this kind of pre-contract contingency can not arise in equilibrium.

The remaining argument then closely follows the partitioned unraveling argument presented for the abridged full-commitment contract. Given that it is implementable, a contract induces the profit-maximizing cooperative outcome as soon as it is negotiated, making deviations most profitable in the period prior to contracting. This leaves players with a partitioned game, where cooperation is induced as soon as the contract comes into effect, and unrestrained play (at best featuring the (l, l) Nash equilibrium in its last period) prevails before. Thus, by this logic, any contract that is written in $t > 0$ is inferior to the baseline. Combined, these two cases give rise to the following result:

Proposition 7. *Every contract chosen in the main model with exogenous timing will also be formed at $t = 0$ under endogenous timing.*

Proposition 7 highlights an aspect that we have taken as a given: Contracts are negotiated at the outset of the game. The proposition demonstrates that, rather than being merely an assumption, early contracting emerges as a property of smooth-landing contracts within our model, even when allowing for other contract timing.

6 Conclusion

Many business relationships rely on loose arrangements in early interactions, only to solidify their alliances through contractual commitments later. More often than not, firms refrain from using rigidly binding agreements from the very beginning of their interactions (Ohmae, 1989). Similarly, firms often offer small and standardized contracts to build rapport with their suppliers and only later move to larger and more complex orders (Bernstein and Peterson, 2022).

We employ a repeated-game framework (with a finite horizon) to investigate how firms can utilize hybrid-contracting strategies—the progression from relational contracts (Baker et al., 2002) in early stages to formal contracts and commitment in later stages—to ensure mutual cooperation throughout the entire business relationship. We demonstrate that smooth-landing contracts, which formalize these end-of-day interactions in a flexible manner, can extend the duration of a cooperative business relationship (at the *intensive margin*) and expand the set of environments in which cooperation can be achieved (at the *extensive margin*). Such contracts extend cooperation more cost-effectively than a full-commitment contract that prescribes the cooperative action in all periods, from the very beginning. Moreover, unlike abridged full-commitment contracts that prescribe the cooperative action only in later periods, smooth-landing contracts do not crowd out early-stage cooperative behavior under the relational contract dynamics.

Our results contribute to the existing economics and management literature on contracts in two major ways, extending beyond the perspective of building and fostering trust, which is often the focus of the contracting literature in practice. First, we provide a justification for the use of hybrid contracting in the contexts of strategic alliances, research joint ventures, and other collaborative arrangements. Our results rationalize contractual solutions intended to overcome alliance failures (e.g., Parkhe, 1993; Hughes and Weiss, 2007) when, initially, loose alliances are successful. The formal, smooth-landing contracts complement, rather than substitute for, informal relational contracts. Our approach, however, is different from, for instance, Baker et al. (2002), by focusing on the dynamics of a shift or progression in the contracting structure (from informal to formal).

Second, the trade-off between wider implementability and higher costs (or lower payoffs) stemming from longer-period smooth-landing contracts aligns with insights from the literature on incomplete contracts. A complete contract—analogous to our full-commitment contract—is effective in implementing the cooperative outcome across a wide range of environments. However, it is costly to write and enforce—more costly than an incomplete contract that does not govern all periods or does not specify a course of action for all contingencies. Instead of writing a full (or complete) contract, parties can agree on a more cost-effective incomplete contract now and fill in the gaps later. In our setting, the relational part is the gap in an incomplete contract, and the smooth-landing contract fills that gap. In our setting, contractual incompleteness (in a hybrid-contracting approach) takes on a temporal dimension.

For practitioners, our paper offers insights into effective contract management. First, formal contracting at the later stages of a business relationship can complement informal arrangements at earlier stages, but must be flexible to avoid crowding out the very relational-contract dynamics it is intended to add to. We demonstrate how a *rigid* abridged full-commitment contract crowds out cooperation, whereas a flexible smooth-landing contract, which restricts the action space rather than prescribing the cooperative action, is more effective. This extends cooperation at later stages while retaining early-stage cooperation. Second, drafting smooth-landing contracts early and enforcing them late prevents crowding out and lowers the effective costs of hybrid contracting.

Third, longer contracts (covering more periods) relax implementability constraints and allow for cooperation across a wider range of circumstances—but these benefits come at a higher cost. Longer contracts can outperform shorter contracts but are optimal only when the longer duration is needed to cover a longer period of non-cooperation (under the relational-contract dynamics). In fact, our results suggest a maximum length of smooth-landing contracts. Longer contracts beyond that upper bound then start to resemble full-commitment contracts that implement the cooperative outcome at unnecessarily high costs.

Lastly, there is a connection between our results and managerial compensation. Time preferences of firms' managers or their executive compensation packages are a critical factor in hybrid contracting. Because time-myopic managers or those with short-term incentive compensation are said to discount future earnings more than current earnings and focus more on short-termism in their decisions (e.g., Stein, 1989; Marinovic and Varas, 2019; Edmans et al., 2022), affected firms need to enter longer smooth-landing contracts. Conversely, firms whose executive compensation structures feature longer-term incentives can rely on shorter and more cost-effective smooth-landing contracts.

References

Anderlini, Luca and Leonardo Felli (1994) "Incomplete Written Contracts: Undescribable States of Nature," *Quarterly Journal of Economics*, 109 (4), 1084–1124.

Andreoni, James and John H. Miller (1993) "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence," *Economic Journal*, 103 (418), 570–585.

Baker, George, Robert Gibbons, and Kevin J. Murphy (2002) "Relational Contracts and the Theory of the Firm," *Quarterly Journal of Economics*, 117 (1), 39–84.

Basov, Suren (2016) *Social Norms, Bounded Rationality, and Optimal Contracts*, vol. 30 of Studies in Economic Theory, Singapore: Springer.

Battigalli, Pierpaolo and Giovanni Maggi (2002) "Rigidity, Discretion, and the Costs of Writing Contracts," *American Economic Review*, 92 (4), 798–817.

——— (2008) "Costly Contracting in a Long-Term Relationship," *RAND Journal of Economics*, 39 (2), 352–377.

Benoit, Jean-Pierre and Vijay Krishna (1985) "Finitely Repeated Games," *Econometrica*, 53 (4), 905–922.

Bernstein, Lisa (2016) "Beyond Relational Contracts: Social Capital and Network Governance in Procurement Contracts," *Journal of Legal Analysis*, 7 (2), 561–621.

Bernstein, Lisa and Brad Peterson (2022) "Managerial Contracting: A Preliminary Study," *Journal of Legal Analysis*, 14 (1), 176–243.

Cubitt, Robin P., Michalis Drouvelis, Simon Gächter, and Ruslan Kabalin (2011) "Moral Judgments in Social Dilemmas: How Bad is Free Riding?," *Journal of Public Economics*, 95 (3-4), 253–264.

Diamond, Douglas W. (1984) "Financial Intermediation and Delegated Monitoring," *Review of Economic Studies*, 51 (3), 393–414.

——— (1991) "Monitoring and Reputation: The Choice between Bank Loans and Directly Placed Debt," *Journal of Political Economy*, 99 (4), 689–721.

Dye, Ronald A. (1985) "Costly Contract Contingencies," *International Economic Review*, 26 (1), 233–250.

Edmans, Alex, Vivian W. Fang, and Allen H. Huang (2022) "The Long-Term Consequences of Short-Term Incentives," *Journal of Accounting Research*, 60 (3), 1007–1046.

Frydlinger, David, Oliver Hart, and Kate Vitasek (2019) "A New Approach to Contracts," *Harvard Business Review*, 97 (5), 116–126.

Fudenberg, Drew and Jean Tirole (1991) *Game Theory*, Cambridge, Mass.: MIT Press.

Gibbons, Robert and Rebecca Henderson (2012) "Relational Contracts and Organizational Capabilities," *Organization Science*, 23 (5), 1350–1364.

Gulati, Ranjay, Tarun Khanna, and Nitin Nohria (1994) "Unilateral Commitments and the Importance of Process in Alliances," *MIT Sloan Management Review*, 35 (3), 61–69.

Gulati, Ranjay and Harbir Singh (1998) "The Architecture of Cooperation: Managing Coordination Costs and Appropriation Concerns in Strategic Alliances," *Administrative Science Quarterly*, 43 (4), 781–814.

Hughes, Jonathan and Jeff Weiss (2007) "Simple Rules for Making Alliances Work," *Harvard Business Review*, 85 (11), 122–128.

Kaplow, Louis (1995) "A Model of the Optimal Complexity of Legal Rules," *Journal of Law, Economics, and Organization*, 11 (1), 150–163.

Klein, Benjamin (2002) "The Role of Incomplete Contracts in Self-Enforcing Relationships," in Brousseau, Eric and Jean-Michel Glachant eds. *The Economics of Contracts: Theories and Applications*, 59–71, Cambridge, UK: Cambridge University Press.

Kreps, David M., Paul Milgrom, John Roberts, and Robert Wilson (1982) "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory*, 27 (2), 245–252.

Liker, Jeffrey K. and Thomas Y. Choi (2004) "Building Deep Supplier Relationships," *Harvard Business Review*, 82 (12), 104–115.

Marinovic, Ivan and Felipe Varas (2019) "CEO Horizon, Optimal Pay Duration, and the Escalation of Short-Termism," *Journal of Finance*, 74 (4), 2011–2053.

Melumad, Nahum, Dilip Mookherjee, and Stefan Reichelstein (1997) "Contract Complexity, Incentives, and the Value of Delegation," *Journal of Economics and Management Strategy*, 6 (2), 257–289.

Moore, Nancy Y., Laura H. Baldwin, Frank Camm, and Cynthia R. Cook (2002) "Implementing Best Purchasing and Supply Management Practices: Lessons from Innovative Commercial Firms," documented briefing, RAND Corporation, Santa Monica, Cal.

Ohmae, Kenichi (1989) "The Global Logic of Strategic Alliances," *Harvard Business Review*, 67 (2), 143–154.

Park, Seung Ho and Michael V. Russo (1996) "When Competition Eclipses Cooperation: An Event History Analysis of Joint Venture Failure," *Management Science*, 42 (6), 875–890.

Park, Seung Ho and Gerardo R. Ungson (2001) "Interfirm Rivalry and Managerial Complexity: A Conceptual Framework of Alliance Failure," *Organization Science*, 12 (1), 37–53.

Parkhe, Arvind (1993) "Strategic Alliance Structuring: A Game Theoretic and Transaction Cost Examination of Interfirm Cooperation," *Academy of Management Journal*, 36, 794–829.

Poppo, Laura and Todd Zenger (2002) "Do Formal Contracts and Relational Governance Function as Substitutes or Complements?," *Strategic Management Journal*, 23, 707–725.

Ryall, Michael D. and Rachelle C. Sampson (2009) "Formal Contracts in the Presence of Relational Enforcement Mechanisms: Evidence from Technology Development Projects," *Management Science*, 55 (6), 906–925.

Schilling, Melissa A. (2019) *Strategic Management of Technological Innovation*, New York, N.Y.: McGraw-Hill Eduction, 5th edition.

Schwartz, Alan and Joel Watson (2004) "The Law and Economics of Costly Contracting," *Journal of Law, Economics, and Organization*, 20 (1), 2–31.

Stein, Jeremy C. (1989) "Efficient Capital Markets, Inefficient Firms: A Model of Myopic Corporate Behavior," *Quarterly Journal of Economics*, 104 (4), 655–669.

Strachan, Stuart (2019) "Cartels and Joint Ventures," *Competition Chronicle*, May 1, 2019, available at http://web.archive.org/web/20200811034651/https://www.competitionchronicle.com/2019/05/cartels-and-joint-ventures/ (last accessed: June 18, 2025).

Tirole, Jean (2009) "Cognition and Incomplete Contracts," *American Economic Review*, 99 (1), 265–294.

Williamson, Oliver E. (1985) *The Economic Institutions of Capitalism*, New York, NY: Free Press.

——— (1989) "Transaction Cost Economics," in Schmalensee, Richard and Robert Willig eds. *Handbook of Industrial Organization*, 1, Chap. 3, 135–182, Amsterdam, NL: North-Holland.

A Appendix

A.1 Proofs of Main Results

Proof of Lemma 1

Proof. We first derive the critical threshold $\bar{\delta}^D$ such that players delay deviations as long as possible whenever $\delta \geq \bar{\delta}^D$. We then show that a partial-cooperation equilibrium is not sustainable for lower δ. Last, we derive the critical threshold for the partial-cooperation equilibrium.

First, for the optimal timing of a deviation assume (w, w) (with payoffs of one) has been played for the first t' periods and compare the payoffs from deviation in $t = t' + 1$ (with payoffs of π^D in $t' + 1$ and 0 in $t' + 2$ and all future periods, if any) with deviation in $t = t' + 2$ (with payoffs of 1 in $t' + 1$ and π^D in $t' + 2$ and 0 in all future periods, if any). Players prefer deviation later if

$$\pi^D + \delta \cdot 0 \leq 1 + \delta\pi^D \qquad \Longleftrightarrow \qquad \delta \geq \frac{\pi^D - 1}{\pi^D} =: \bar{\delta}^D.$$

By induction, it follows that whenever $\delta < \bar{\delta}^D$, it is optimal to deviate in $t = 1$, whereas for $\delta > \bar{\delta}^D$, deviating in $t = T - 1$ is the most profitable.

Second, to show that a partial-cooperation equilibrium is not sustainable for $\delta < \bar{\delta}^D$, suppose for a moment that it is. Because for $\delta < \bar{\delta}^D$, a player deviates (if it does) in $t = 1$ (with payoffs π^D in $t = 1$ and 0 in all future periods), cooperation in $t = 1$ is profitable (and suppose cooperation in all future periods except in $t = T$) if

$$\frac{1 - \delta^T}{1 - \delta} + \delta^T \pi^L \geq \pi^D + \delta \cdot 0 + \dots$$

The present discounted value of an infinite income stream $\pi_t = 1$ for infinite periods is $\frac{1}{1-\delta}$. It is easy to see that these payoffs are higher than receiving $\pi_t = 1$ for T periods and $\pi_t = \pi^L < 1$ in the $(T + 1)^{\text{th}}$ (i.e., the last). Moreover, we have $\frac{1}{1-\delta} \leq \pi^D$ if, and only if $\delta \leq \bar{\delta}^D$. Collecting terms, we obtain

$$\pi^D \geq \frac{1}{1 - \delta} > \frac{1 - \delta^T}{1 - \delta} + \delta^T \pi^L,$$

contradicting the working assumption that cooperation (payoffs on the RHS) is more profitable than deviation (payoffs on the LHS). For $\delta < \bar{\delta}^D$, a partial-cooperative equilibrium is, therefore, not sustainable.

Third, in the last period, the outcome is the Nash equilibrium (l, l) with payoffs π^L. A player will cooperate if, in $T - 1$, the payoffs from cooperation are higher than from deviation (the latter triggering payoffs of 0 from the Nash equilibrium (s,s) in T), or

$$1+\delta\pi^L \geq \pi^D \quad \Longleftrightarrow \quad \delta \geq \frac{\pi^D-1}{\pi^L} =: \bar{\delta}^{BL}.$$

Given cooperation in $t=T-1$, a player will cooperate in $T-2$ if

$$1+\delta+\delta^2\pi^L \geq \pi^D \quad \Longleftrightarrow \quad \delta\left(1+\delta\right) \geq \frac{\pi^D-1}{\pi^L};$$

a player will cooperate in $T-3$ if

$$1+\delta+\delta^2+\delta^3\pi^L \geq \pi^D \quad \Longleftrightarrow \quad \delta\left(1+\delta\left(1+\delta\right)\right) \geq \frac{\pi^D-1}{\pi^L};$$

and so forth. Hence, if $\delta \geq \bar{\delta}^{BL}$, a player cooperates in $T-1$ and all preceding periods.

Because $\pi^D > \pi^L$ we have $\bar{\delta}^D < \bar{\delta}^{BL}$, finalizing the proof. □

Proof of Lemma 2

Proof. The proof is trivial as playing a Nash equilibrium in every subgame is, by definition, a subgame-perfect Nash equilibrium. □

Proof of Proposition 1

Proof. The proof is by the discussion in the main text and the formal derivations of the end-of-game dynamics in Lemma 1. □

Proof of Lemma 3

Proof. If in $T-1$, players cooperate, the strategy implies an outcome (w,w) in T (with payoffs of 1). If a player deviates, the strategy implies an outcome (l,l) in T (with payoffs of π^L). A player will cooperate in $T-1$ if the payoffs from cooperation are higher than from deviation, or

$$1+\delta \geq \pi^D+\delta\pi^L \quad \Longleftrightarrow \quad \delta \geq \frac{\pi^D-1}{1-\pi^L} =: \bar{\delta}^{SL}.$$

Given cooperation in $t=T-1$, a player will cooperate in $T-2$ if

$$1+\delta+\delta^2 \geq \pi^D+\delta\pi^L+\delta^2\pi^L \quad \Longleftrightarrow \quad \delta\left(1+\delta\right) \geq \frac{\pi^D-1}{1-\pi^L};$$

and so forth. Hence, if $\delta \geq \bar{\delta}^{SL}$, a player cooperates in $T-1$ (triggering cooperation in T) and all preceding periods. Because $\pi^D > 1 > \pi^L$ (immediately implying $\pi^D > 1-\pi^L$) we have $\bar{\delta}^D < \bar{\delta}^{SL}$, finalizing the proof. □

Proof of Lemma 4

Proof. If in $T-1$, players cooperate, the strategy implies an outcome (w, w) in T (with payoffs of 1). If a player deviates, the strategy implies an outcome (s, s) in T (with payoffs of 0). A player will cooperate in $T-1$ if the payoffs from cooperation are higher than from deviation, or

$$1 + \delta \geq \pi^D + \delta \cdot 0 \quad \iff \quad \delta \geq \pi^D - 1 =: \bar{\delta}^{SL^+}.$$

Given cooperation in $t = T-1$, a player will cooperate in $T-2$ if

$$1 + \delta + \delta^2 \geq \pi^D \delta \cdot 0 + \delta^2 \cdot 0 \quad \iff \quad \delta\,(1+\delta) \geq \pi^D - 1;$$

and so forth. Hence, if $\delta \geq \bar{\delta}^{SL^+}$, a player cooperates in $T-1$ (triggering cooperation in T) and all preceding periods. Because $\pi^D > 1$ we have $\bar{\delta}^D < \bar{\delta}^{SL^+}$, finalizing the proof. □

Proofs of Propositions 2 and 3

Proof. The derivations for the cost thresholds in comparison to the partial-cooperative and the non-cooperative baseline for the full-commitment contract, as well as the simple and improved smooth-landing contracts, are each given in the corresponding sections. □

Proof of Proposition 4

Proof. Proposition 4 collects the results from Lemmas 1–4 and Propositions 2 and 3.

1. The first point follows from only full-commitment contracts being implementable (Lemma 3 and 4) and profitable if κ lies below the relevant cost threshold (Proposition 2). The alternative baseline equilibrium is given by the implementability threshold of the baseline (Lemma 1).

2. The second point follows from the implementability threshold of the improved smooth-landing contract (Lemma 4) together with the fact that the cost threshold for a smooth-landing contract (Proposition 3) is lower than that of a full-commitment contract (Proposition 2), implying lower costs and, thus, payoff dominance. The alternative baseline equilibrium is given by the implementability threshold of the baseline (Lemma 1).

3. The third point follows from the implementability threshold of both smooth-landing contracts (Lemma 3 and 4) together with the fact that the cost threshold for both smooth-landing contracts (Proposition 3) is lower than that of a full-commitment contract (Proposition 2), implying lower costs and, thus, payoff dominance. The

alternative baseline equilibrium is given by the implementability threshold of the baseline (Lemma 1).

4. The fourth point follows from the implementability threshold of the improved smooth-landing contract (Lemma 4) together with the fact that the cost threshold for a smooth-landing contract (Proposition 3) is lower than that of a full-commitment contract (Proposition 2), implying lower costs and, thus, payoff dominance. The alternative baseline equilibrium is given by the implementability threshold of the baseline (Lemma 1).

5. The last point follows from the implementability threshold of both smooth-landing contracts (Lemma 3 and 4) together with the fact that the cost threshold for both smooth-landing contracts (Proposition 3) is lower than that of a full-commitment contract (Proposition 2), implying lower costs and, thus, payoff dominance. The alternative baseline equilibrium is given by the implementability threshold of the baseline (Lemma 1). □

Proof of Proposition 5

Proof. To show that any smooth-landing contract is (jointly) optimal whenever implementable and profitable, first, we show that the joint first-best outcome is not achievable absent any commitment device; second, that the smooth-landing contracts are among the most cost-efficient class of contracts that achieve the first-best if implementable; and lastly, that this commitment-dependent first-best is preferable to the second-best whenever the smooth-landing contracts are profitable.

First, note that the joint stage-game payoff is maximized by (w, w). Hence, the joint first-best outcome is the fully cooperative outcome of (w, w) being played in all periods. To show that this outcome is not achievable without commitment, assume to the contrary the existence of a set of strategies $\hat{\alpha}$ implementing this outcome. It is trivial to see that deviating to $\hat{\alpha}_{i,T} = s$ is a profitable deviation for any player, as $\pi^D > 1$. Thus, no set of strategies can exist that implements the fully cooperative outcome. This logic extends to any implemented outcome that does not feature a Nash equilibrium in the last stage. Hence, the second-best has to feature the payoff-dominant (l, l) Nash equilibrium in the last period. Note that the baseline grim trigger strategy (Strategy 2.2), implements this outcome in the last period, and the most profitable outcome in any period prior, and hence, the second-best outcome without commitment.

Second, given the first-best is not achievable without commitment, doing so requires the usage of contracts (as it is the only available commitment device in our framework). Due to the cost structure of contracting, commitment comes at a lump-sum cost of κ. Thus, the final cost of commitment depends solely on the timing of the cost realization.

With smooth-landing contracts enforced only in the last period, this immediately implies that they are part of the cost-minimizing class of contracts. Further, if implementable, they are able to implement the first-best fully cooperative outcome.

Lastly, for the class of cost-minimizing, first-best implementing contracts (including implementable smooth-landing contracts) to be optimal, they need to payoff-dominate the costless second-best. Because payoffs of all contracts in said class are necessarily identical, this condition is analogous to the profitability condition of the smooth-landing contracts (Proposition 3), concluding the proof. □

A.2 Formal Derivations and Proofs of Extensions

A.2.1 Multi-Period Smooth-Landing Contracts

Proof of Lemma 5

Proof. To show that $\bar{\delta}^k(\tau)$ is decreasing in τ for any smooth-landing contract type k, assume for now that players prefer to deviate as late as possible ($\delta \geq \bar{\delta}^D$). Fix an arbitrary length of contract τ and denote by $\pi_k^{\text{Pun}} \in \left\{\pi^L, 0\right\}$ the payoff each player receives after deviating. Thus, cooperation in $T - \tau$ (the period before the contract is enforced) is optimal if

$$1 + \delta + \ldots + \delta^\tau \geq \pi^D + \left(\delta + \ldots + \delta^\tau\right) \pi_k^{\text{Pun}} \quad \iff \quad \delta + \ldots + \delta^\tau \geq \frac{\pi^D - 1}{1 - \pi_k^{\text{Pun}}}.$$

Rearrange and consider a strict equality (i.e., the player is indifferent between cooperation and deviation), then $\left(\delta + \ldots + \delta^\tau\right)\left[1 - \pi_k^{\text{Pun}}\right] = \pi^D - 1$. Define $\tilde{\delta}_k(\delta') = \delta' + \ldots + \delta'^\tau$ where δ' is the discount factor such that the above equality holds. Because both $1 - \pi_k^{\text{Pun}}$ and $\pi^D - 1$ do not change with δ and τ, $\tilde{\delta}_k(\delta')$ must not change with τ. As $\tilde{\delta}_k(\delta')$ is increasing in δ', we must have a lower δ' for higher values of τ. This implies that the critical threshold $\bar{\delta}^k(\tau)$ is decreasing in τ.

By the reasoning in Lemma 1, we cannot obtain cooperation for $\delta < \bar{\delta}^D$, and the implementability threshold $\bar{\delta}^k(\tau)$ is bound below by $\bar{\delta}^D$. If T is big enough, there is a $\tau' =: \bar{\tau}^k$ such that $\bar{\delta}^k(\tau') \geq \bar{\delta}^D$ and τ'' such that $\bar{\delta}^k(\tau'') < \bar{\delta}^D$. □

Proof of Proposition 6

Proof. The proof follows from Lemma 5 and the increase in effective contracting costs as τ increases. □

A.2.2 Contracting with Upfront Costs

By introducing a more general cost structure with drafting and enforcement costs, the profitability thresholds stated in Propositions 2 and 3 change. Below, we derive the modified thresholds in Table 4.

Smooth-Landing Contract for Partial-Cooperative Baseline. For the smooth-landing contracts to be profitable, we require

$$\begin{aligned} \frac{1-\delta^T}{1-\delta} + \delta^T(1-\kappa_i) - \kappa_d = \frac{1-\delta^T}{1-\delta} + \delta^T(1-\alpha\kappa) - (1-\alpha)\kappa &\geq \frac{1-\delta^T}{1-\delta} + \delta^T \pi^L \\ \iff \kappa\left(\delta^T\alpha + 1 - \alpha\right) &\leq \delta^T\left(1-\pi^L\right) \\ \iff \kappa &\leq \frac{\delta^T}{1-\alpha+\alpha\delta^T}\left(1-\pi^L\right) \end{aligned}$$

when compared with the partial-cooperative baseline.

Smooth-Landing Contract for Non-Cooperative Baseline. For the smooth-landing contracts to be profitable, we require

$$\begin{aligned} \frac{1-\delta^{T+1}}{1-\delta} - \delta^T\kappa_i - \kappa_d = \frac{1-\delta^{T+1}}{1-\delta} - \delta^T\alpha\kappa - (1-\alpha)\kappa &\geq \frac{1-\delta^{T+1}}{1-\delta}\pi^L \\ \iff \kappa(\delta^T\alpha + 1 - \alpha) &\leq \frac{1-\delta^{T+1}}{1-\delta}\left(1-\pi^L\right) \\ \iff \kappa &\leq \frac{1-\delta^{T+1}}{(1-\delta)(1-\alpha+\alpha\delta^T)}\left(1-\pi^L\right) \end{aligned}$$

when compared with the non-cooperative baseline.

Full-Commitment Contract. Because, for the full-commitment contract, all contracting costs are paid in the initial period, the distinction into drafting and enforcement costs has no effect, and the thresholds described in Proposition 2 remain unchanged.

A.2.3 Endogenous Contract Timing

Proof of Proposition 7

Proof. As in Section 5.2, assume that costs are a combination of drafting costs $\kappa_d = (1-\alpha)\kappa$ and enforcement costs $\kappa_i = \alpha\kappa$.

Consider *first* the scenario in which only the non-cooperative baseline exists ($\delta < \bar{\delta}^{BL}$). We compare an implementable smooth-landing contract agreed upon in period τ with a contract agreed upon in $\tau + 1$. Note that, as discussed in Lemmas 3 and 4, implementability does not depend on the length of the game—or the length of the respective subgames in this case. For the player to prefer the τ contract, we require

$$\delta^{\tau}(1-\kappa_d)+\delta^{\tau+1}>\delta^{\tau}\pi^L+\delta^{\tau+1}(1-\kappa_d)$$
$$\iff \kappa_d<\frac{1-\pi^L}{1-\delta}.$$

This threshold is independent of τ and, hence, by induction, any smooth-landing contract is either optimal in $t=0$ or $t=T$.

Now consider a smooth-landing contract optimal in the last stage and compare it to the non-cooperative baseline. The latter is optimal whenever $\delta^T\pi^L>\delta^T(1-\kappa)$, which is equivalent to $\kappa>1-\pi^L$. This last condition is necessarily true because $\kappa>\kappa_d>\frac{1-\pi^L}{1-\delta}>1-\pi^L$. Hence, either a smooth-landing contract is negotiated at $t=0$ or no smooth-landing contract is formed, which simplifies to the standard model of Section 3.

In the *second scenario*, we have $\delta\geq\bar{\delta}^{BL}$, implying the existence of a partial-cooperative baseline equilibrium. We need to consider the possibility of unilateral deviations in an intermediate period, hence, we approach the game using backward induction. In the last period, following a non-deviation history, the payoff without contract is π^L while after a deviation history it is 0. Under an existing contract, the payoff is $1-\kappa$.[23] Players prefer to form a contract in all histories if $\kappa<1-\pi^L$ and prefer to form a contract only in deviation histories (but not in others) if $\kappa<1$. Because $\kappa<1-\pi^L$ is required for smooth-landing contracts to be profitable compared to the partial-cooperative baseline irrespective of game length (see Proposition 3), we focus on the former case with $\kappa<1-\pi^L$.

As discussed previously, contracts must be independent of prior play histories to be sequentially rational. In that case, the last-period contract yields $1-\kappa$ irrespective of whether players deviated previously. In period $T-1$, playing (w,w) is no longer optimal since players would deviate and still enter a non-punishing contract afterwards. Hence, players in a no-deviation history face the choice of sticking to the second-best π^L or of forming a contract that yields $1-(1-\delta)\,k_d$. In both cases, players earn $1-\kappa$ in the last period. For not contracting in a given period, we need $\pi^L>1-(1-\delta)\,\kappa_d$, which is equivalent to $\kappa_d>\frac{1-\pi^L}{1-\delta}$ but contradicts $\kappa_d\leq\kappa<1-\pi^L$. Hence, it is optimal to contract in that period. The same must hold for players in a deviation history facing a zero payoff without a contract. Hence, again, after both kinds of histories, a contract would be entered that partitions the game in the same way as before. Crucially now, in each preceding history, the highest attainable payoff without a contract is π^L after no-deviation histories and 0 after deviation histories. The last step is to show that in any period, a contract adds $1-(1-\delta)\,\kappa_d$. We show this by complete induction. Start from a last-period contract in which the total payoff is $1-\kappa$. Now take an arbitrary contract

[23] Technically, it is only $1-\kappa_i$ if the contract was drafted in earlier periods, but this assumption simplifies analysis and will be accounted for in later stages.

that yields $1-\kappa_d+\delta+\ldots+\delta^{n-1}+\delta^n(1-\kappa_i)$ and note that we obtain the previous period total value of $1-\kappa_d+\delta+\ldots+\delta^n+\delta^{n+1}(1-\kappa_i)$ by adding $1-(1-\delta)\kappa_d$ to the discounted value, i.e.,

$$\begin{aligned} & 1-(1-\delta)\kappa_d+\delta\left(1-\kappa_d+\delta+\ldots+\delta^{n-1}+\delta^n(1-\kappa_i)\right) \\ =\ & 1-\kappa_d+\delta\kappa_d+\delta-\delta\kappa_d+\ldots+\delta^n+\delta^{n+1}(1-\kappa_i) \\ =\ & 1-\kappa_d+\delta+\ldots+\delta^n+\delta^{n+1}(1-\kappa_i) \end{aligned}$$

Thus, the game transforms for each subsequent period into a stationary game where players decide between π^L and $1-(1-\delta)\kappa_d$ in no-deviation histories and 0 and $1-(1-\delta)\kappa_d$ in deviation histories, followed by a safe contract irrespective of their choice in that period. Thus, the only smooth-landing contracts that arise are when $\kappa<1-\pi^L$ and are agreed upon at $t=0$. □

A.3 Mixed-Strategy Equilibria of the Stage Game

Existence of Mixed-Strategy Equilibria. To derive mixed-strategy Nash equilibria (MNE) of the stage game, first note that pure-strategy action w is dominated by s:

$$\begin{aligned} &1) \quad 0>2(\pi-c)=2\left(\tfrac{1}{2}-\pi\right) \text{ by Condition 1(a)} \\ &2) \quad 2\pi-m>4\pi-2c=1 \\ &\quad \iff m<2\pi-1 \text{ by Condition 1(c)} \\ &3) \quad \tilde{e}\pi-m>(2+\tilde{e})\pi-2c=(2+\tilde{e})\pi-4\pi+1 \\ &\quad \iff m<2\pi-1 \text{ by Condition 1(c)} \end{aligned}$$

As a consequence, the only MNE is with the mixed strategy $\alpha_i=ps+(1-p)l$. The indifference condition for MNE then yields $1-p=\frac{\tilde{e}}{m}\left(\pi-\frac{1}{2}\right)$. Because $\pi_i(s,l)<\tilde{e}/2$ and $\pi_i(l,s)<\tilde{e}/2$, it follows that the expected payoffs from the MNE is below π^L. For $\pi_i(s,l)$ this is satisfied because (l,l) is a Nash equilibrium, for $\pi_i(l,s)$ this follows from the fact that $\pi_i(l,s)=\tilde{e}(\pi-c)<\tilde{e}(2\pi-c)=\frac{1}{2}\tilde{e}$ because $\tilde{e}>0$ and $\pi>0$. Finally, the modified condition for (d), that is, $\frac{m}{\pi}<\tilde{e}<\min\left\{\frac{m}{\pi-1/2},\frac{m+1/2}{\pi}\right\}$, is obtained from the condition ensuring $\tilde{e}\pi-m>0$, which in turn implies a positive payoff from the MNE.

Without this assumption, the MNE could provide a negative payoff and offer an alternative to punishing deviations (other than (s,s)). In this case, let $\pi^{\text{MNE}}:=p\cdot 0+(1-p)(\tilde{e}\pi-m)=-\frac{(\tilde{e}\pi-m)\tilde{e}(1/2-\pi)}{m}$ denote the stage-payoffs of the MNE. Note that this is saying $\pi^{\text{MNE}}=-\frac{\pi_i(s,l)\pi_i(l,s)}{m}$ which is another way to obtain the modified condition for (d) because $\pi_i(l,s)<0$ by Condition 1. With this additional Nash equilibrium of the stage game, the grim trigger strategy used in the baseline can be refined using the MNE as punishment (instead of (s,s)), leading to a new threshold $\bar{\delta}^{D'}=\frac{\pi^D-1}{\pi^L-\pi^{\text{MNE}}}$.

Using Mixed-Strategy Equilibria in Contracting. Similar to using $\pi_i(s, s)$ as a more effective punishment in the context of improved smooth-landing contracts in Section 3.2.2, the MNE can be used as a punishment to further improve the ability of the smooth-landing contract to implement full cooperation. This would affect the δ-thresholds with a corresponding new implementability threshold of $\bar{\delta}^{\text{MNE}} = \frac{\pi^D - 1}{1 - \pi^{\text{MNE}}}$, which, for the reasons discussed above, satisfies $\bar{\delta}^{\text{MNE}} < \bar{\delta}^{SL^+} < \bar{\delta}^{SL}$. While extending the smooth-landing contract's area of implementability even further, the usage of such MNE in the grim trigger strategy, however, does not change the subsequent results qualitatively.

Download ZEW Discussion Papers:

https://www.zew.de/en/publications/zew-discussion-papers

or see:

https://www.ssrn.com/link/ZEW-Ctr-Euro-Econ-Research.html
https://ideas.repec.org/s/zbw/zewdip.html

//
IMPRINT

ZEW – Leibniz-Zentrum für Europäische Wirtschaftsforschung GmbH Mannheim

ZEW – Leibniz Centre for European Economic Research

L 7,1 · 68161 Mannheim · Germany
Phone +49 621 1235-01
info@zew.de · zew.de