



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Dynamic relational contracts under complete information

Citation for published version:

Thomas, J & Worrall, T 2018, 'Dynamic relational contracts under complete information' *Journal of Economic Theory*, vol. 175, pp. 624-651. DOI: 10.1016/j.jet.2018.02.004

Digital Object Identifier (DOI):

[10.1016/j.jet.2018.02.004](https://doi.org/10.1016/j.jet.2018.02.004)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Journal of Economic Theory

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.





Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 175 (2018) 624–651

JOURNAL OF
**Economic
Theory**

www.elsevier.com/locate/jet

Dynamic relational contracts under complete information

Jonathan P. Thomas, Tim Worrall *

School of Economics, University of Edinburgh, Edinburgh, United Kingdom

Received 12 February 2016; final version received 3 February 2018; accepted 7 February 2018

Available online 12 February 2018

Abstract

This paper considers a long-term relationship between two agents who both undertake an action or investment that produces a joint benefit. Agents have an opportunity to expropriate some of the joint benefit for their own use. Agents have quasi-linear preferences. Two cases are considered: where agents are risk averse but where limited liability constraints do not bind, and where agents are risk neutral and subject to limited liability constraints. We ask how to structure the investments and division of the surplus over time to avoid expropriation. In the risk-averse case, the dynamics of actions and surplus may or may not be monotonic depending on whether or not a first-best allocation can be sustained. Agents may underinvest but never overinvest. If the first-best allocation is not sustainable, there is a trade-off between risk sharing and surplus maximization; surplus may not be at its constrained maximum even in the long run and the “amnesia” property of pure risk-sharing models fails to hold. In contrast, in the risk-neutral case there may be an initial phase in which one agent overinvests and the other underinvests. Both actions and surplus converge monotonically to a stationary state, where surplus is maximized subject to the self-enforcing constraints.

© 2018 Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

JEL classification: C61; C73; D86; D91; L14

Keywords: Relational contracts; Self-enforcement; Limited commitment; Risk sharing

* Corresponding author.

E-mail addresses: jonathan.thomas@ed.ac.uk (J.P. Thomas), tim.worrall@ed.ac.uk (T. Worrall).

1. Introduction

This paper considers a situation where two agents repeatedly engage in joint production. In each period, both agents simultaneously undertake an action or investment that produces a joint output. Agents must also decide how to share the joint output each period. We assume there is a hold-up problem, that is, contracts on actions or the division of the joint output are not enforceable and in addition the outside option of each agent is increasing in the investment of the other agent. We allow joint output and the outside options of the agents to depend on an exogenous state. We consider cases where the agents are risk averse and where they are risk neutral. The only link between periods is a Markov process determining states. There is complete information: apart from the fact that the agents choose their actions simultaneously each period, everything is observable. The only friction is that contracts cannot be enforced. We consider allocations or contracts from which no agent has an incentive to renege by imposing self-enforcing constraints at each date and state. We refer to feasible contracts that satisfy these constraints as dynamic relational contracts. We characterize the Pareto-efficient dynamic relational contracts; we refer to such contracts as optimal contracts.

We impose two simplifying assumptions on our model. First, we assume that agents' preferences are quasi-linear in consumption and actions. This simplifies the problem because with quasi-linear preferences efficient actions (and hence, surplus) are determined independently of the distribution of resources (the marginal rate of substitution between consumption and the action is equal to unity). Second, we impose sufficient conditions such that the constrained Pareto-frontier is concave. This simplifies our problem because it allows us to focus on non-random contracts.¹ We examine two main cases: where agents are risk averse but preferences are such that non-negativity constraints on consumption can be ignored, and where agents are risk neutral but consumption is constrained to be non-negative (limited liability).

If agents are risk averse results depend on whether or not it is possible to sustain a first-best allocation for some division of the surplus. If it is possible, convergence to the first best is monotone. Otherwise there might be an initial monotone phase, but in the long-run, when there are two or more states, monotonicity does not generally obtain: when the same state recurs, surplus will sometimes be higher at the later date and sometimes lower. There is also a trade-off between achieving efficient risk-sharing and maximizing current surplus even in the long run. In particular, and in contrast to the risk-neutral case, current surplus is not maximized. Better risk-sharing is achieved by holding the action of one agent inefficiently low because this reduces the outside option of the other agent, that is, it relaxes the latter's self-enforcing constraint. We show that the optimal contract depends on the past history of states and so the "amnesia" property of the risk-sharing limited commitment model does not hold.

When agents are risk neutral, we consider the implications of limited liability constraints and show that optimal contracts involve two phases. In the first phase there is backloading with zero consumption for the constrained agent, who overinvests up to the last period of the backloading phase and the terms of the contract move monotonically in his/her favor. This overinvestment arises because it allows a further transfer of utility to the other agent who consumes the extra output. It occurs despite the hold-up problem, that in a static model would lead to underinvestment. Nevertheless, we demonstrate that because of backloading it is never the case that both

¹ It would be straightforward to allow for random contracts by introducing a public randomization device, but at the cost of considerable complexity of notation and statements of our results. Furthermore, the assumptions we make are consistent with those that are commonly made in the literature.

agents overinvest—even at different dates—in any optimal contract. The second phase is stationary and independent of the initial conditions. Consumption and investment depend on the state but not on the time period. Each agent has positive consumption and, for a given state, either both invest efficiently or both underinvest. In either case, current surplus is maximized subject to the self-enforcing constraints. Convergence to the stationary phase is monotone in the sense that whenever the same state recurs in the backloading phase, surplus is higher at the later date.

1.1. Related literature

A number of results for special or limiting cases of this model are known. First, one-sided-action versions of this model or variations on it, have been studied by a number of authors (see, e.g., Thomas and Worrall, 1994; Sigouin, 2003; Albuquerque and Hopenhayn, 2004; Kovrijnykh, 2013). Typically, this literature has considered the case where both agents are risk neutral, there is limited liability and the agent taking the action can commit. To prevent the uncommitted agent from taking his/her outside option, actions may be kept low initially. A key insight of this literature is that incentives are improved when payments to the uncommitted agent are backloaded into the future. This provides a growing carrot for adhering to the contract. Consequently, the action or investment of the other agent can be increased in the future. This generates dynamics in the agent's actions as well as in monetary payments. In the long run, actions and transfers converge to a stationary distribution that maximizes the surplus, output less action costs, given the self-enforcing constraints. The speed of backloading is restricted by the limited liability constraints. Ray (2002) has established the most general backloading result of this type. He considers a general, but non-stochastic, principal-agent model in which both parties may take actions. The principal can commit within each period, so the self-enforcing constraint only applies to the agent. He shows that an efficient contract has terms that move in favor of the agent, converging in finite time to the efficient self-enforcing continuation that maximizes the agent's payoff. Our results generalize this backloading result to the case where both agents undertake an action and neither agent can commit. Furthermore, we demonstrate that there may be overinvestment in the risk-neutral case, and in the risk-averse case show that there may be a trade-off between productive efficiency and risk sharing even in the long run. Neither of these properties occur in models where only one agent takes an action.

Second, consider the case where agents have no action to take, or where there is no hold-up problem. In this case, the model involves sharing a stochastic endowment. The case in which agents have their own stochastic endowment and can share risk subject to limited commitment constraints has been widely studied (see, e.g., Kocherlakota, 1996; Ligon et al., 2002; Thomas and Worrall, 1988). A result of this pure risk-sharing case is that a constrained Pareto-efficient allocation evolves toward a stationary distribution, and that, for some parameter values, the distribution of future expected utilities is non-degenerate. Although the distribution is non-degenerate, the solution exhibits an “amnesia” property that once an agent is constrained, the contract from then on is independent of the past history of shocks. With hold-up, the optimal contract depends on the past history of states and does not in general exhibit the amnesia property of the pure risk-sharing model.² Furthermore, the pure risk-sharing literature only considers distributional issues and has no implications for the efficiency and dynamics of actions that are the focus of

² Ábrahám and Lacsó (2018) establish a similar result in a model of risk-sharing model and storage. The absence of the amnesia property is more consistent with the empirical evidence (see Broer, 2013).

this paper. Nevertheless, we are able to demonstrate a limit result that as our hold-up problem vanishes, the optimal contract converges to the standard pure risk sharing contract.

Third, there are a very few papers in this limited commitment literature that examine the situation where two or more agents take actions. The most relevant paper to ours is Acemoğlu et al. (2011) that considers a model of changes in political power. In Acemoğlu et al. (2011) a Markov process determines which risk-averse political party is in power. Political parties take actions that contribute to a common pool of resources whether in power or not, but only the party in power gets to determine the allocation of resources across agents. Therefore, states are identified by the agent in power. It is shown that in a constrained Pareto-efficient allocation, the action of one of the agents (the one in power) is always chosen efficiently and actions of other agents (those not in power) are distorted downward. Furthermore, they establish a convergence result that depends on whether a first-best allocation is sustainable or not: if a first-best allocation is sustainable, then the actions and the division of resources converges to a degenerate (first-best) distribution; otherwise, allocations converge to a non-degenerate distribution (it need not be unique). The two-agent model with quasi-linear utility considered in their paper corresponds to the limiting case of our model where in each state one agent has all the property rights. We also establish convergence results but our results apply for a general distribution of property rights and an arbitrary number of states and may result in the actions of both agents being inefficiently low, even in the long run. Their convergence result, when a first-best allocation is sustainable, corresponds to our Theorem 1(a). In Theorem 1(b), when a first-best allocation is not sustainable, we establish convergence to a unique limiting distribution that is independent of initial conditions.

Fourth, our model is related to the broader literature on relational contracting (see, e.g., Levin, 2003; Doornik, 2006; Rayo, 2007) that builds on the work of Macleod and Malcolmson (1989). This literature has studied models with more general ingredients (including many-sided actions, enforceable payments, moral hazard, hidden information, and endogenous property rights), but has restricted attention to stationary equilibria, thus, eliminating any interesting dynamics in investments and transfers. The restriction to stationary equilibria is either derived, because stationary contracts are optimal (when agents are risk neutral and in the absence of limited liability), or imposed, because the focus is on organizational structures under which full efficiency can be achieved. Most of this literature is therefore silent on the dynamics of relational contracts that are the main concern of this paper.³

1.2. Illustrative example

To illustrate the model we have in mind, we present a simple example with no uncertainty and risk averse agents.⁴ There are two agents with common discount factor $\delta \in (0, 1)$ and an infinite horizon. In each discrete period the action or effort of agent i is a_i and joint output is additive

³ One exception to the focus on stationary contracts is Fong and Li (2017) who introduce limited liability and moral hazard into a risk-neutral model firms and workers based on Levin (2003). They show that if the principal extracts most of the surplus, the backloading of the agent's utility can lead to a probationary contract in which the agent's wage is initially at the lower bound, and incentives are provided by the threat of termination; at some point this threat is removed and the wage increases to a higher level.

⁴ For the purpose of constructing a simple example that illustrates the solution, we here ignore the non-negativity constraints on consumption. We use parameter values such that the Pareto-frontier is concave. In the Supplementary Material, we show how to fully solve this example using our characterization results and without having to use value function iteration.

$$y(a_1, a_2) = f_1(a_1) + f_2(a_2) = 2(\sqrt{a_1} + \sqrt{a_2}).$$

Both agents have common preferences satisfying constant absolute risk aversion with coefficient 1/2:

$$u_i(x) = 2\left(1 - e^{-\frac{1}{2}x}\right),$$

where $x := c - a$, consumption less effort. Actions take place simultaneously at the beginning of each period. At the end of the period, output is realized and it is divided between the two agents. Suppose, that irrespective of how output is divided, agent i can unilaterally get a breakdown consumption of $\phi_i(a_1, a_2) = \theta_{i1}f_1(a_1) + \theta_{i2}f_2(a_2)$, that depends on the action of the other agent. For parameters $\theta_{11} = \theta_{22} = 0$ and $\theta_{12} = \theta_{21} = 1$, this means that either agent can expropriate all of the other agent’s output but if they do so they lose their own output. A relational contract is just an agreed sequence of actions and division of the output from these actions from which no agent has an incentive to deviate. We assume that if a deviation occurs, in each period thereafter the agents revert to short-run Nash equilibrium anticipating the breakdown payoffs $\phi_i(a_1, a_2)$. With the specification for θ_{ij} just given, the short-run Nash equilibrium has $a_i, c_i = 0$ (and hence, $u_i = 0$) and the discounted payoff from a deviation, the deviation utility, is:

$$D_i(a_j) = u_i(2\sqrt{a_j}) = 2\left(1 - e^{-\sqrt{a_j}}\right).$$

We characterize *constrained* Pareto optimal contracts, that is, within the set from which no agent would deviate. At the first best, $a_1^* = a_2^* = 1$ and surplus $z := y(a_1^*, a_2^*) - a_1^* - a_2^* = 2$ is maximal. This is sustainable provided an equal split of surplus ($x_i^* = 1$) is an equilibrium:

$$\frac{u_i(1)}{(1 - \delta)} = \frac{2(1 - e^{-\frac{1}{2}})}{(1 - \delta)} \geq D_i(1) = 2(1 - e^{-1}),$$

or $\delta \geq (1 + \sqrt{e})^{-1}$.

Suppose that $\delta \geq (1 + \sqrt{e})^{-1}$, so that the first-best allocation is sustainable. Let V_i denote the lifetime utility of agent i . For $\delta > (1 + \sqrt{e})^{-1}$ surplus will be constant at its efficient level for a range of values for V_1 . Consider starting from a feasible value of V_1 below $D_1(1) = u_1(2)$, i.e., worse for agent 1 than the deviation utility at the first-best allocation. If $a_2 = a_2^*$, then agent 1 would deviate. Therefore, $a_2 < a_2^*$; the best contract has a_2 as high as possible, such that agent 1 does not wish to deviate, i.e., $V_1 = D_1(a_2)$. Since we are assuming $\delta \geq (1 + \sqrt{e})^{-1}$, for the corresponding value of V_2 on the Pareto frontier, $V_2 > u_2(2)$. That is, agent 2 is unconstrained and $a_1 = 1$ (is efficient). Since a_2 , and hence, surplus z , is determined by the binding constraint $V_1 = D_1(a_2)$, both can be expressed as functions of V_1 . Hence, in this example:

$$\begin{aligned} a_1(V_1) &= 1, \\ a_2(V_1) &= \left(\frac{1}{2}u_1^{-1}(V_1)\right)^2 = \left(\log\left(1 - \frac{V_1}{2}\right)\right)^2, \\ z(V_1) &= 1 + u_1^{-1}(V_1) - \left(\frac{1}{2}u_1^{-1}(V_1)\right)^2 = 1 - 2\log\left(1 - \frac{V_1}{2}\right) - \left(\log\left(1 - \frac{V_1}{2}\right)\right)^2, \\ z'(V_1) &= u_1^{-1'}(V_1)\left(1 + u_1^{-1}(V_1)\right) = \left(1 + \log\left(1 - \frac{V_1}{2}\right)\right)\left(1 - \frac{V_1}{2}\right)^{-1}. \end{aligned}$$

It is easily checked that $z(V_1)$ is increasing and concave in this region with $z(0) = z'(0) = 1$ and $z(u_1(2)) = 2$ and $z'(u_1(2)) = 0$. We show below (equation (4.1) in Section 4) that for values of

V_1 where the surplus is increasing in V_1 , as here for $V_1 < u_1(2)$, then V_1 will be higher next period. It follows straightforwardly that V_1 is an increasing sequence converging to $u_1(2)$. So, the contract converges to the surplus maximizing actions, here the first best. We also show below (see equation (3.2b) in Section 3) that surplus is divided so that u'_2/u'_1 is equal to the absolute value of the slope of the (strictly concave) Pareto frontier in the following period. Since V_1 is increasing over time, so too is u'_2/u'_1 .⁵ That is, the way the surplus is distributed (as well as V_1) moves monotonically in favor of agent 1. Thus, backloading of agent 1's utility occurs, and in such a way as to guarantee efficiency in the long-run.

The case where $\delta < (1 + \sqrt{e})^{-1}$ is similar. There is convergence to a stationary value of V_1 . However, in this case, convergence of V_1 is to a point where the actions maximize the joint surplus subject to the no-deviation constraints. At this value of V_1 , and for some neighborhood around it, both constraints $V_i = D_i(a_j)$ bind.

Taking both cases together, it can be concluded that with no uncertainty there is convergence to the constrained surplus maximizing actions for any δ .

1.3. Plan of paper

The paper extends this example to consider a more general production function and breakdown payoffs. We consider multiple states and quasi-linear preferences including risk-neutrality and non-negativity constraints on consumption. We show that the convergence result of the example generalizes to the case with multiple states when the first-best is sustainable (and also when agents are risk-neutral), but otherwise, with risk aversion and multiple states there is a trade-off between risk sharing and efficiency and convergence to surplus maximization does not occur. In this case we show that u'_2/u'_1 converges to a non-degenerate limiting distribution independent of the initial distribution of the surplus.

The paper proceeds as follows. Section 2 describes the model. Section 3 provides some general results that apply to both the risk-neutral and risk-averse cases. Section 4 analyzes the risk-averse case and Section 5 the risk-neutral case. Section 6 concludes. Statements of lemmas and the proofs of theorems are found in Appendix A. Proofs of Propositions and Lemmas are relegated to the Supplementary Material.

2. Model

We consider a dynamic model of joint production where agents repeatedly undertake an action or investment that generates a joint output. There is no asset accumulation and full depreciation of the investment in each period. Once produced agents have the opportunity to unilaterally expropriate some of the joint output for their own benefit. In this section, we shall describe the economic environment and the set of dynamic relational contracts. We define a game played by the two agents and identify dynamic relational contracts with the subgame perfect equilibria of that game. Our interest is in optimal contracts, that correspond to the set of Pareto-efficient subgame perfect equilibria.⁶

⁵ Convergence of u'_2/u'_1 is to $e(1 - \delta(1 - e))^{-2} \leq 1$; convergence is to 1 for $\delta = (1 + \sqrt{e})^{-1}$.

⁶ More precisely, we focus on efficient pure subgame-perfect equilibria relative to specified "Nash reversion" punishments, although our characterization also applies mutatis mutandis to optimal punishments, should they be different, and hence, to efficient equilibria among the set of all pure strategy equilibria.

2.1. Economic environment

Time is discrete and indexed by $t = 0, 1, 2, \dots, \infty$. At the start of each period, a state of nature s is realized from a finite state space \mathcal{S} with $n \geq 1$ states. The state evolves according to an irreducible, time homogeneous Markov chain with transition matrix $[\pi_{sr}]$, where $\sum_{r \in \mathcal{S}} \pi_{sr} = 1$, all $s \in \mathcal{S}$. The chain starts from an initial state s_0 at date $t = 0$. We denote the state at date t by s_t and the history of states by $s^t = \{s_0, s_1, \dots, s_t\}$.

There are two agents, $i = 1, 2$. At every date t , and after the state at that date is observed, both agents simultaneously choose an action/investment $a_i \in \mathbb{R}_+$. Actions produce an output $y^s(a) \geq 0$ that depends on the state s and the action pair $a := (a_1, a_2)$ (details are given below in Assumption 2). Having observed actions and output, the agents agree to split output and each consumes non-negative consumption $c_i, c := (c_1, c_2) \in \mathbb{R}_+^2$. We impose that consumption is non-negative as a simple way to reflect a limited liability constraint on the transfers one agent can make to the other. Consumption c is *feasible* if $c_1 + c_2 \leq y^s(a)$. Agent i derives per-period utility u_i from *net consumption* $x_i := c_i - a_i, x := (x_1, x_2) \in \mathbb{R}^2$. We make the following assumptions on u_i and y^s :

Assumption 1. Per-period utility $u_i: [\underline{x}_i, \infty) \rightarrow \{-\infty\} \cup \mathbb{R}$ is a twice continuously differentiable, strictly increasing and concave function of net consumption, where $\underline{x}_i \leq 0$.

Assumption 2. For each $s \in \mathcal{S}$, the production function $y^s: \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$ is twice continuously differentiable, strictly increasing in both arguments and strictly concave. Furthermore, for each $s \in \mathcal{S}, \partial^2 y^s(a) / \partial a_1 \partial a_2 \geq 0$ (complementarity); $y^s(0) = 0$ and the upper contour sets $\{a \in \mathbb{R}_+^2 \mid y^s(a) - a_1 - a_2 \geq \gamma\}, \gamma \in \mathbb{R}$, are compact.

Assumption 2 imposes fairly standard conditions on the production function. The last part of Assumption 2 is a simple way to restrict actions to a compact set $\mathcal{A}(s)$. Denote *surplus* in state s by $z^s(a) := y^s(a) - a_1 - a_2$. Define the *first-best action pair* $a^*(s)$ as the actions that maximize surplus in state s . Given Assumption 2, the first-best action pair exists and is unique. We refer to the surplus $z^s(a^*(s))$ as the *first-best surplus*. Since actions are chosen simultaneously and independently, we also define the *conditionally efficient actions* $a_i^*(a_j, s), i, j = 1, 2, i \neq j$, such that

$$a_i^*(a_j, s) := \arg \max_{a_i \in \mathbb{R}_+} [y^s(a_1, a_2) - a_i].$$

The conditionally efficient actions are single-valued, continuous functions of the other agent’s action.⁷ The weak complementarity assumption is slightly restrictive but reflects our view that relational contracting framework is most natural when there are complementarities in production. Given the weak complementarity assumption, conditionally efficient action functions are weakly upward sloping. In addition, $a_i^*(s) = a_i^*(a_j^*(s), s)$ for $i, j = 1, 2, i \neq j$.

We now specify what an agent can get if there is no agreement on how to divide up output. If no agreement is reached, agent i gets a *breakdown consumption* of $\phi_i^s(a)$, and hence, a *breakdown utility* of $u_i(\phi_i^s(a) - a_i)$. An agent can always take the option of receiving her breakdown

⁷ This result is simple and straightforward to show. Formal statements of results of this type are stated as lemmas in Appendix A and proofs are given in the Supplementary Material. For example, the fact that conditionally efficient actions are single-valued and continuous is stated as Lemma 1 in Appendix A.

utility. More formally, we suppose the agents play a Nash demand game to divide output.⁸ In this Nash demand game, both agents simultaneously announce consumption claims $(\tilde{c}_1, \tilde{c}_2)$, $\tilde{c}_i \geq 0$. If $\tilde{c}_1 + \tilde{c}_2 = y^s(a)$, then this determines the division of output: consumption $c_i = \tilde{c}_i$. Otherwise, agents receive their breakdown consumption: $c_i = \phi_i^s(a)$.

The specific assumptions on $\phi_i^s(a)$ are given below, but a simple example with *proportional defaults* captures what we have in mind. Suppose that each agent can, by defaulting, capture a fraction θ_i of the available output $y^s(a)$. Here, $\phi_i^s(a) = \theta_i y^s(a)$. We assume that agents cannot obtain more than the available output, so $\theta_1 + \theta_2 \leq 1$. We do not require that the sum exhausts available output. For example, disagreement may incur a cost, such as lawyers' fees or bargaining costs, so that some of output is lost when there is default. In such cases, $\theta_1 + \theta_2 < 1$. We assume $\theta_i > 0$, so that what an agent gets in the breakdown is increasing in the action of the other agent. This assumption captures the hold-up feature of joint production we wish to model.

As another example, consider the special case with *additive* production: $y^s(a) = f_1^s(a_1) + f_2^s(a_2)$ and suppose $\phi_i^s(a) = \theta_{i1}^s f_1^s(a_1) + \theta_{i2}^s f_2^s(a_2)$, $\theta_{ij}^s \geq 0$ and $\sum_{i=1}^2 \theta_{ij}^s \leq 1$, $j = 1, 2$ (this is very similar to the formulation used by Halonen (2002)). Our hold-up assumption requires $\theta_{ij}^s > 0$, $i, j = 1, 2$, $i \neq j$. With this parameterization, assuming $\sum_{i=1}^2 \theta_{ij}^s = 1$ and taking the limit as $\theta_{ij}^s \rightarrow 0$, for $i, j = 1, 2$, $i \neq j$ and for all $s \in \mathcal{S}$, produces the pure risk sharing model that has been studied by Kocherlakota (1996); Ligon et al. (2002) and others. This is discussed in Section 4.

Analogous to Assumption 2, we shall assume that $\phi_i^s(a)$ satisfies:

Assumption 3. For each $s \in \mathcal{S}$ and $i = 1, 2$, the function $\phi_i^s: \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$ is continuous, twice continuously differentiable, strictly increasing in both arguments and strictly concave. Moreover, $\partial^2 \phi_i^s(a) / \partial a_1 \partial a_2 \geq 0$ (complementarity) and $\partial \phi_i^s(0, a_j) / \partial a_i > 1$ for all $a_j \in \mathbb{R}_+$, $i, j = 1, 2$, $i \neq j$. In addition, $\phi_i^s(0, 0) = 0$ for $i = 1, 2$ and

$$\frac{\partial \phi_1^s(a)}{\partial a_1} + \frac{\partial \phi_2^s(a)}{\partial a_1} \leq \frac{\partial y^s(a)}{\partial a_1} \quad \forall s \text{ and } i = 1, 2. \tag{2.1}$$

In the case of proportional defaults, these conditions (apart from $\partial \phi_i^s(0, a_j) / \partial a_i > 1$) follow directly from Assumption 2. Complementarity in Assumption 3 implies that the reaction functions in the breakdown game are weakly upward sloping, and this simplifies the arguments below. Condition (2.1) requires that the marginal change in the total breakdown consumption from a change in the action of one of the agents cannot exceed the corresponding marginal product. Together with $\phi_i^s(0, 0) = 0$, it implies that the $\phi_i^s(a)$ are feasible, that is, $\phi_1^s(a) + \phi_2^s(a) \leq y^s(a)$ for each a and s . Condition (2.1) together with $\partial \phi_i^s(0, a_j) / \partial a_i > 1$, $i = 1, 2$ implies that the first-best action pair is strictly positive. The assumption that ϕ_i^s is strictly increasing in both its arguments, in particular that $\partial \phi_i^s(a) / \partial a_j > 0$ for $i \neq j$, captures the hold-up property of the model.

Denote the Nash best-response functions (functions because $\phi_i^s(a_1, a_2)$ is strictly concave in a_i) in the breakdown game by

$$a_i^N(a_j, s) := \arg \max_{a_i \in \mathbb{R}_+} [\phi_i^s(a_i, a_j) - a_i].$$

⁸ This is approach is also used by Hall (2005), for example. What we want to capture is that there is an ex ante agreement on what actions should be taken, and how the resulting output should be split, and that failure to abide by it leads to the breakdown utilities. The Nash demand game is a simple way of implementing this idea – but we stress that our results are not sensitive to the way it is operationalized. For a fuller discussion, see Hall (2005).

The Nash best response function $a_i^N(a_j, s)$ is continuous and weakly increasing in a_j . Moreover, we have $0 < a_i^N(a_j, s) < a_i^*(a_j, s)$ for each a_j and every state $s \in \mathcal{S}$. It is strictly positive because $\partial\phi_i^s(0, a_j)/\partial a_i > 1$ and is less than the conditionally efficient action because of the hold-up assumption that $\partial\phi_i^s(a)/\partial a_j > 0$. The best-response breakdown utility is

$$u_i^N(a_j, s) := u_i(\phi_i^s(a_i^N(a_j, s), a_j) - a_i^N(a_j, s)).$$

A Nash equilibrium of the breakdown game occurs where the best-response functions intersect (existence follows by standard arguments). Without further assumptions, the Nash equilibrium need not be unique (though it is unique if the defaults are proportional). However, the potential non-uniqueness is not critical because the Nash equilibria can be Pareto-ranked (because the best-response functions are non-decreasing and all Nash equilibria lie below the first-best action pair $a^*(s)$). Henceforth, we let $(a_1^{NE}(s), a_2^{NE}(s))$ denote the dominant Nash equilibrium and all our results apply relative to this dominant Nash equilibrium.

2.2. Dynamic relational contracts

We refer to a non-negative action and consumption sequence $\{a(s^t), c(s^t)\}_{t \geq 0}$ as a *plan*. Corresponding to a plan, agent i 's lifetime utility is

$$V_i(s_0) := \mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t u_i(c_i(s^t) - a_i(s^t)) \mid s_0 \right],$$

where δ is a common discount factor, $0 < \delta < 1$, and \mathbb{E} denotes expectation. A plan is feasible if $\sum_i c_i(s^t) \leq y^{s^t}(a(s^t))$ for every history s^t and $c_i(s^t) - a_i(s^t) \geq \underline{x}_i$ for $i = 1, 2$ and every history s^t .

A *dynamic relational contract*, or simply *contract*, is a feasible plan from which neither agent has an incentive to deviate. The incentive to deviate depends on the punishment for deviation. This is given by the breakdown payoffs in the current period (subsequent to the deviation), and by play of the (dominant) equilibrium of the static breakdown game in all future periods. In particular, suppose that a is the current recommended action pair. If agent i is to deviate at t , then the best she can do is to choose $a_i^N(a_j(s^t), s_t)$, which yields a current payoff $u_i^N(a_j(s^t), s_t)$.⁹ She is punished from $t + 1$ by “Nash reversion” in which both agents choose their best responses in the breakdown game, that is, both will thereafter play the (dominant) Nash equilibrium of the breakdown game described above.¹⁰

Let $D_i^s(a_j)$ denote the *deviation utility*: the best discounted payoff that agent i can get by deviating, given agent j 's putative action a_j in state s . It is defined recursively by

$$D_i^s(a_j) := u_i^N(a_j, s) + \delta \sum_{r \in \mathcal{S}} \pi_{sr} D_i^r(a_j^{NE}(r)),$$

where $D_i^r(a_j^{NE}(r))$ is the deviation utility from the play of the Nash equilibrium in state r . Given our hold-up assumption (see Assumption 3), it follows that the deviation utility is continuous, differentiable, strictly increasing and strictly concave in the action of the other agent.

⁹ Deviation at the output division stage cannot be preferable since breakdown is triggered in either case, and a_i may not be optimal in the breakdown.

¹⁰ A dynamic relational contract is equivalent to a pure strategy subgame perfect equilibrium relative to future reversion to this Nash equilibrium. Here, strategies are infinite sequences of history-dependent actions and consumption *claims*. Punishment consisting of immediate triggering of the breakdown, and repeated play of the (dominant) Nash equilibrium of the breakdown game thereafter, is subgame perfect (each agent just *demand*s the whole output after any deviation (i.e., off the equilibrium path), triggering the breakdown game each period).

We stress that replacing the Nash reversion punishments by any state dependent continuation utilities that are no greater than the Nash reversion punishments leaves all the characterization results we derive intact. In particular, optimal punishments satisfy this property. Equally, if agents can take state-dependent outside options at the start of any period, then, provided these outside options satisfy the condition that they are no greater than the Nash reversion punishments, all our results apply. For example, if in periods after a default the breakdown consumptions/utilities were lower than they are in an on-going relationship, then our results still hold.

Since an agent can always take the option of receiving her breakdown utility, the deviation utility provides a lower bound (as a function of the other agent’s action) on the discounted utility an agent gets in any dynamic relational contract. Hence, $\{a(s^t), c(s^t)\}_{t=0}^\infty$ is a dynamic relational contract if it is feasible and if for every s^t , and $i, j = 1, 2, i \neq j$,

$$V_i(s^t) := u_i(c_i(s^t) - a_i(s^t)) + \mathbb{E} \left[\sum_{\tau=t+1}^\infty \delta^{\tau-t} u_i(c_i(s^\tau) - a_i(s^\tau)) \mid s^t \right] \geq D_i^{s^t}(a_j(s^t)). \tag{2.2}$$

The *continuation utility* $V_i(s^t)$ is the discounted utility that agent i anticipates from the contract after the history s^t . The right hand side of (2.2) is the deviation utility agent i gets from deviating from the recommended action after the history s^t . We refer to the inequalities (2.2) as the *self-enforcing* constraints. Whenever (2.2) holds with equality, we say that agent i is *constrained*. Otherwise, we say that agent i is *unconstrained*.

Dynamic relational contracts exist. For example, the *trivial* contract that has $a_i(s^t) = a_i^{NE}(s_t)$ and $c_i(s^t) = \phi_i^{s^t}(a^{NE}(s_t))$ for all s^t is both feasible and self-enforcing and therefore a dynamic relational contract. We show below (see Proposition 2) that there exist other non-trivial dynamic relational contracts.¹¹ Corresponding to any dynamic relational contract, $\{a(s^t), c(s^t)\}_{t=0}^\infty$, and initial state s_0 , is a pair of lifetime utilities $(V_1(s_0), V_2(s_0))$. Given the set of dynamic relational contracts, let \mathcal{V}_{s_0} denote the set of the corresponding lifetime utilities. Our objective is to characterize contracts corresponding to the Pareto-frontier of the set \mathcal{V}_{s_0} . We refer to dynamic relational contracts that correspond to this Pareto-frontier as *optimal contracts* and refer to the corresponding actions as *optimal actions*. We say that agent i *underinvests* (or that the action is *inefficiently low*) at some date t in an optimal contract if the optimal actions are such that $a_i(s^t) < a_i^*(a_j, s)$ and say the agent *overinvests* (or the action is *inefficiently high*) if $a_i(s^t) > a_i^*(a_j, s)$. Given the stochastic history s^t , we can treat an optimal contract as a stochastic process for (a, c) . We will be interested in the long-run behavior of this process and whether it converges, and if so, whether convergence is dependent on s_0 or $V_1(s_0)$.

3. Preliminary results

This section establishes some preliminary results on the Pareto-frontier of the set of dynamic relational contracts and optimal actions. Section 4 considers the case where agents are risk averse and Section 5 will consider the case where agents are risk neutral.

¹¹ Intuitively, hold-up creates an inefficiency and provided $\delta > 0$, repeated game arguments allow cooperation to improve on the breakdown Nash equilibrium.

3.1. Relationship to the Nash actions

Proposition 1. *In any optimal contract (i) actions are never below the Nash reaction functions, $a_i(s^t) \geq a_i^N(a_j(s^t), s_t)$, and $a(s^t) \geq a^{NE}(s_t) > 0$ for all s^t ; (ii) an agent who is allocated all current output and who is not overinvesting (i.e., $a_i(s^t) \leq a_i^*(a_j(s^t), s_t)$), is unconstrained.*

The intuition for (i) is that if the action of one of the agents, say agent 1, were below the Nash reaction function, a Pareto improvement could be found by increasing the action of agent 1 by a small amount. Although the deviation utility of agent 2 increases (by hold-up), his consumption can be increased to prevent a violation of his self-enforcing constraint, and there is sufficient extra output remaining to more than compensate agent 1 for the increase in her action. This property then implies that actions can never be below the Nash equilibrium actions, $a(s^t) \geq a^{NE}(s_t)$. Since it can be shown that the Nash equilibrium actions are strictly positive, $a_i^N(a_j, s) > 0$, it follows that optimal actions are always positive too. Although (ii) is not trivial, it is unsurprising. Suppose, say, that agent 1 is allocated all of the current output. Then, agent 1 is receiving more of output than she would obtain in the breakdown game, if she held her action constant (because, by Assumption 3, agent 2 can claim a positive share of output in the breakdown game). In a deviation, agent 1 will optimize her action, but since she is not overinvesting, reducing her action to the Nash reaction function will only reduce output net of her effort. Hence, she would be worse off than receiving all output at the higher action. The continuation utility cannot be lower than the deviation continuation utility, so a deviation will lead to output being shared and a punishment continuation, worse than the equilibrium path and thus, agent 1 could not be constrained. In fact, we shall show later (see the discussion after the first-order conditions (3.2a)–(3.2c)) that any agent with positive consumption will not overinvest, and therefore the caveat in Proposition 1(ii) about an agent who is not overinvesting can be dispensed with.

3.2. Concavity, continuity and differentiability

We define $V_2^s(V_1)$ to be the Pareto-frontier of the set \mathcal{V}_s . It is not necessarily concave; in particular the concavity of the deviation utility $D_j(a_i)$ in the action of the other agent implies that the self-enforcing constraints (2.2) may not be satisfied at average actions and hence the constraint set need not be convex. Nevertheless, the Pareto-frontier can be shown to be concave under some additional restrictions. We state and discuss two *alternative* sufficient conditions for concavity in Appendix A, Assumption A4 and Assumption A5.

Proposition 2. *For each $s \in \mathcal{S}$ (i) under either Assumption A4 or Assumption A5, $V_2^s(V_1)$ is a continuous and concave function of V_1 defined on a non-degenerate closed interval $[\underline{V}_1^s, \bar{V}_1^s]$, and is continuously differentiable on its interior. Moreover,*

$$V_2^{s(+)}(\underline{V}_1) = 0 \quad \text{and} \quad V_2^{s(-)}(\bar{V}_1) = -\infty,$$

where $V_2^{s(+)}$ denotes the right and $V_2^{s(-)}$ the left derivative. (ii) Under Assumption A4, $V_2^s(V_1)$ is strictly concave if u_i is strictly concave, $i = 1, 2$, or over any interval such that $a^s(V_1)$ varies with V_1 ; under Assumption A5, $V_2^s(V_1)$ is strictly concave over any interval such that $a^s(V_1)$ varies with V_1 .

We use Assumption A4 in Section 4 that considers the case where agents are risk averse. It requires two things: the first is essentially that the curvature of the deviation utility is less than the

curvature of surplus as a function of actions. The second is that an optimal contract has $x_i > 0$, $i = 1, 2$ at every date. The latter follows, for example, for utility functions (such as those with constant relative risk aversion with coefficient of risk aversion greater than or equal to one) where $\lim_{x \rightarrow 0} u(x) = -\infty$. We use Assumption A5 in Section 5 that considers the case where agents are risk neutral. It requires that the production function is more concave than the corresponding deviation utility. It is satisfied in many reasonable examples and Assumption A5 is a generalized version of the condition given in Thomas and Worrall (1994).

In the one-sided action case where only one agent undertakes an action, it is known that the value function can fail to be differentiable (Thomas and Worrall, 1994). It is perhaps surprising, then, that in this two-sided case we are able to establish differentiability. The key observation is that since optimal actions are positive, it is possible to vary both actions simultaneously, holding the future utilities constant, so as to vary V_1 whilst satisfying the self-enforcing and feasibility constraints.

3.3. Recursive formulation

We now use a recursive programming approach to examine optimal contracts. It is useful to work with net consumption x_i as a choice variable instead of consumption c_i . The Markov assumption on the evolution of states and the infinite time horizon, together with the observation that all the self-enforcing constraints are forward looking, means that the set of continuation utilities corresponding to a dynamic relational contract depends only on the state r and is independent of the past history. $V_2^s(V_1)$ is characterized as follows:

Proposition 3. $V_2^s(V_1)$, $V_1 \in [V_1^s, \bar{V}_1^s]$, is a solution to the following program

$$V_2^s(V_1) = \max_{a \geq 0, x \geq \underline{x}, (V_1^r \in \mathbb{R})_{r \in \mathcal{S}}} \left\{ u_2(x_2) + \delta \sum_{r \in \mathcal{S}} \pi_{sr} V_2^r(V_1^r) \right\} \tag{P1}$$

subject to

$$u_1(x_1) + \delta \sum_{r \in \mathcal{S}} \pi_{sr} V_1^r \geq V_1: \quad \lambda \tag{3.1a}$$

$$V_1 \geq D_1^s(a_2): \quad \mu_1 \tag{3.1b}$$

$$u_2(x_2) + \delta \sum_{r \in \mathcal{S}} \pi_{sr} V_2^r(V_1^r) \geq D_2^s(a_1): \quad \mu_2 \tag{3.1c}$$

$$V_1^r \geq \underline{V}_1^r: \quad \delta \pi_{sr} \nu_1^r \tag{3.1d}$$

$$V_1^r \leq \bar{V}_1^r: \quad \delta \pi_{sr} \nu_2^r \tag{3.1e}$$

$$x_i + a_i \geq 0 \quad i, j = 1, 2, \quad i \neq j \quad \gamma_i \tag{3.1f}$$

$$x_1 + x_2 \leq z^s(a_1, a_2): \quad \psi \tag{3.1g}$$

The non-negative Lagrangian multipliers are indicated after each inequality. The expected discounted utility V_1 of agent 1 (in state s) is the state variable in this programming problem. The value function $V_2^s(V_1)$ represents the Pareto-frontier of the set of dynamic relational contracts in the space of continuation utilities. It describes how the maximum continuation utility to agent 2 changes as the continuation utility of agent 1 is changed. The inequality (3.1a) is the *promise-keeping constraint* that requires that the contract delivers at least the current discounted utility. The inequalities (3.1b) and (3.1c) are the self-enforcing constraints corresponding to the

inequalities given in (2.2). The constraints (3.1d) and (3.1e) reflect that the continuation utility for agent 1 in state r must lie in the interval $[\underline{V}_1^r, \bar{V}_1^r]$. Inequalities (3.1f) and (3.1g) are the feasibility constraints.

We denote a solution to [P1] by $(a^s(V_1), x^s(V_1))$ and continuation utilities $(V_1^{s,r}(V_1))$. It can be shown that $a^s(V_1)$ is unique; however, $x^s(V_1)$ and $V_1^{s,r}(V_1)$ need not be. Corresponding to this solution, and abusing notation, we define the surplus $z^s(V_1) := z^s(a_1^s(V_1), a_2^s(V_1))$. We discuss the properties of $z^s(V_1)$ below, but we refer to the maximal value of $z^s(V_1)$ for $V_1 \in [\underline{V}_1^s, \bar{V}_1^s]$ as the *constrained maximal surplus* and the actions that maximize this surplus as the *constrained surplus-maximizing* (CSM) actions. Let $\bar{a}(s)$ denote the CSM action in state s .¹² If the CSM actions are equal to the first-best actions $\bar{a}(s) = a^*(s)$ (and hence the constrained maximal surplus equals the first-best surplus), then we say that the first-best is *sustainable* in state s . We denote the set of states in which the first-best actions are sustainable as $\mathcal{S}_* \subseteq \mathcal{S}$ and denote its complement by \mathcal{S}_*^c (it is possible that $\mathcal{S}_* = \emptyset$ or $\mathcal{S}_*^c = \emptyset$). A first-best allocation (FBA) will involve the first best actions, $a^*(s)$, in each state and date *and* complete risk-sharing (that is, net consumption $x^*(s)$ with $x_1^*(s) + x_2^*(s) = z^s(a^*(s))$ such that $u_2'(x_2^*(s))/u_1'(x_1^*(s))$ is constant over all states and dates).

An optimal contract is computed recursively. Start from some given initial value for agent 1’s lifetime utility, $V_1(s_0)$ in state s_0 . The solution to the programming problem provides optimal values for $a(s_0)$ and $x(s_0)$ in state s_0 by setting $V_1 = V_1(s_0)$ in [P1]. The solution also determines the continuation utilities for $V_1^{s_0,r}(V_1(s_0))$ in each possible subsequent state r . At date $t = 1$ and history $s^1 = (s_0, s_1)$, the value for V_1 is determined by the solution for the continuation utility at date $t = 0$ for the appropriate state and the solution to the date $t = 1$ programme determines $a(s^1)$ and $x(s^1)$. The process is repeated to determine $\{a(s^t), x(s^t)\}_{t=0}^\infty$. Doing this for each $V_1(s_0) \in [\underline{V}_1^{s_0}, \bar{V}_1^{s_0}]$ determines the set of optimal contracts.

3.4. First-order conditions

From Proposition 2 the Pareto-frontier is continuously differentiable and the range of absolute slopes of the frontier is $\mathbb{R}_+ \cup \{\infty\}$. Let $\sigma_s(V_1) := -V_2^{s'}(V_1)$ and $\sigma_{s,r}^+(V_1) := -V_2^{r'}(V_1^{s,r}(V_1))$ be the (absolute) slopes of the Pareto-frontiers, where $\sigma_s: [\underline{V}_1^s, \bar{V}_1^s] \rightarrow \mathbb{R}_+ \cup \{\infty\}$ is strictly increasing since the Pareto-frontier is strictly concave. The envelope condition for [P1] is $-\sigma_s(V_1) = -\lambda + \mu_1$. Using this condition, differentiating with respect to x_i, a_i and V^r in [P1], and rearranging gives the first-order conditions:

$$\sigma_{s,r}^+(V_1) - \sigma_s(V_1) = -\sigma_s(V_1) \frac{\mu_2}{1 + \mu_2} + \frac{\mu_1}{1 + \mu_2} + \frac{v_1^r - v_2^r}{1 + \mu_2} \tag{3.2a}$$

$$\sigma_{s,r}^+(V_1) = \frac{u_2'(\cdot)}{u_1'(\cdot)} + \frac{\gamma_2 - \gamma_1}{u_1'(\cdot)(1 + \mu_2)} + \frac{v_1^r - v_2^r}{1 + \mu_2} \tag{3.2b}$$

$$\frac{\mu_j}{1 + \mu_2} \frac{dD_j^s}{da_i} = \frac{\partial z^s}{\partial a_i} \left(u_2'(\cdot) + \frac{\gamma_2}{1 + \mu_2} \right) + \frac{\gamma_i}{1 + \mu_2} \quad i, j = 1, 2 \text{ and } i \neq j. \tag{3.2c}$$

¹² In principle, there may be dynamic relational contracts in which there are actions that achieve a higher surplus but at the cost of lower future surplus. Our definition considers only optimal contracts. However, in both the risk-neutral and risk-averse cases that we consider below, the two concepts coincide and the CSM actions do maximize $z^s(d_1, d_2)$ subject to the self-enforcing constraints. It will also be shown below that in the cases we consider, the CSM actions are unique.

Since the range of absolute slopes of the frontier is $\mathbb{R}_+ \cup \{\infty\}$, it is intuitive that $\sigma_{s,r}^+(V_1)$ is the same for each future state $r \in \mathcal{S}$. To see this first suppose that $v_1^r > 0$. In this case $V_1^{s,r}(V_1) = \underline{V}_1^r$, $\sigma_{s,r}^+(V_1) = 0$ and by a complementary slackness condition $v_2^r = 0$. Then using equation (3.2a), $-\sigma_s(V_1) - \mu_1 = v_1^r > 0$ which gives a contradiction since $\sigma_s(V_1)$ and μ_1 are non-negative. A similar argument can be made to show that $v_2^r = 0$. Since $v_i^r = 0$, it follows from (3.2a) that $\sigma_{s,r}^+(V_1)$ is independent of r and we write $\sigma_s^+(V_1)$ for this common future value. This property greatly simplifies the dynamics of the contracting problem.

It follows directly from the first-order conditions (3.2c) that in an optimal contract (i) there is only ever underinvestment, $a_i(s^t) < a_i^*(a_j(s^t), s)$, if at least one of the agents is constrained; and (ii) if agent i has positive consumption, then he/she does not overinvest, $a_i(s^t) \leq a_i^*(a_j(s^t), s)$. To see the intuition for the first part, suppose that agent 1 is unconstrained. If agent 2 were underinvesting, he could increase investment and generate more surplus. The surplus would be enough to compensate him for the extra investment and agent 1 won't default because she is unconstrained. Thus, it would be possible to find a better contract, yielding a contradiction. Similarly, to see the second part, suppose that agent 1 is overinvesting. Then she could reduce her investment. This relaxes agent 2's self-enforcing constraint (keeping consumptions now and future promises the same). However, output has fallen, so aggregate consumption must fall. If agent 1 has positive consumption, it is possible to keep the consumption of agent 2 the same, while the utility of agent 1 increases because she has cut her investment from above the conditionally efficient level.

There is also a simple corollary to these results: a) both agents cannot be overinvesting (because one agent must have positive consumption); b) an agent cannot be permanently overinvesting because consumption must be positive at some future date – otherwise the self-enforcing constraint would not be satisfied.

4. Risk aversion

For this section we assume that agents are risk averse: we strengthen Assumption 1 and assume that u_i is strictly concave for $i = 1, 2$, and use Assumption A4 from Appendix A. In particular, it is assumed that net consumption and hence, consumption is strictly positive in an optimal contract. It will follow from this that overinvestment is not a feature of an optimal contract. The allocation of net consumption between agents may vary, potentially considerably, across states even in the long-run. Thus, it is important to examine how allowing for risk aversion affects optimal contracts.

4.1. Characterization of optimal contracts

In this sub-section, we consider some properties of the optimal contract and surplus as V_1 varies in a given state, and how the contract is updated period-by-period: in particular, how the ratio of marginal utilities changes from one period to the next. In the following sub-section, we consider the long-run properties of the optimal contract showing that it evolves towards a stationary distribution and study when this stationary distribution does or does not depend on the value of agent 1's lifetime utility $V_1(s_0)$.

Proposition 4. *With risk-averse agents and under Assumption A4 (i) there is no overinvestment, $\partial z^s(a(s^t))/\partial a_i \geq 0$, $i = 1, 2$, all s^t ; (ii) surplus $z^s(V_1)$ is a concave differentiable function (strictly concave if $s \in \mathcal{S}_*^c$) with maximum at unique CSM actions; (iii) at V_1 such that $z^s(V_1)$*




Fig. 1. Surplus Function $z^s(V_1)$.

is maximized, if $s \in S_*^c$ both constraints bind and $a^s(V_1) < a^*(s)$, and if $s \in S_*$ efficient actions $a^*(s)$ are sustainable by definition; (iv) for each $V_1 \in [V_1^s, \bar{V}_1^s]$, $\sigma_s^+(V_1)$, the (absolute value of the) common slope of the Pareto-frontiers next period, and $\sigma_s(V_1)$, the slope of the current Pareto-frontier, satisfy

$$\sigma_s^+(V_1) - \sigma_s(V_1) = u'_2(x_2^s(V_1)) \frac{dz^s(V_1)}{dV_1}. \tag{4.1}$$

The intuition for (i) was discussed above in Section 3.4 when $c_i > 0$ for $i = 1, 2$. Properties (ii) and (iii) are illustrated in Fig. 1.¹³ Equation (4.1) in part (iv) is fundamental to understanding the dynamics of an optimal contract. It is easy to interpret. Consider a (small) unit increase in V_1 . The effect on agent 2’s discounted utility is to change it by approximately $V_2^s'(V_1) = -\sigma_s(V_1)$ units. One way to effect this change (as good as any other at the optimum) is to hold the current utility of agent 1 constant (giving any change in the current surplus to agent 2) and increase V_1^r in each state r , the next-period continuation utilities of agent 1, by $1/\delta$. The effect on agent 2’s current utility is $u'_2(x_2^s(V_1)) \times (dz^s(V_1)/dV_1)$. The effect on the discounted continuation utility of agent 2 is to decrease it by $\sigma_s^+(V_1)$, the same for all future states. The combined effect for agent 2 is $u'_2(x_2^s(V_1)) \times (dz^s(V_1)/dV_1) - \sigma_s^+(V_1)$. Since the overall change in utility for agent 2 is $-\sigma_s(V_1)$, we can equate to get equation (4.1).

The implication for the dynamics of optimal contracts is illustrated in Fig. 1. Consider starting from a value of V_1 below the level that maximizes surplus. In this region, agent 1’s constraint binds ($D_1^s(a_2^s(V_1)) = V_1$) and a_2 is kept inefficiently low ($a_2^s(V_1) < a_2^*(a_1^s(V_1), s)$) to prevent agent 1 from deviating. In this region, V_2 may be high enough to allow a_1 to be conditionally efficient ($a_1^s(V_1) = a_1^*(a_2^s(V_1), s)$) without violating agent 2’s constraint, but if $s \in S_*^c$, then, closer

¹³ Where $\bar{\chi}_1^s$ and $\underline{\chi}_1^s$ are the values for V_1 such that agent 1’s constraint binds for $V_1 \leq \bar{\chi}_1^s$, while agent 2’s constraint binds for $V_1 \geq \underline{\chi}_1^s$; surplus is maximized at $\hat{\chi}_1^s$ in case (b). See the Supplementary Material for further details. Note that both constraints bind for values of $V_1 \in (\underline{\chi}_1^s, \bar{\chi}_1^s)$ in Fig. 1b. This contrasts with pure risk-sharing models with limited commitment, for example, Kocherlakota (1996) or Thomas and Worrall (1988), where at most one self-enforcing constraint binds at any one time in any non-trivial optimum.

to the surplus maximizing value of V_1 , both constraints will bind and a_1 will be inefficiently low. Also, in this region, $d\zeta^s(V_1)/dV_1 > 0$, so equation (4.1) implies that $\sigma_s^+(V_1) > \sigma_s(V_1)$. In particular, if there is a single state or if the same state recurs, the change in V_1 is as indicated by the arrows in Fig. 1. In this case, surplus will be higher next period as the increase in V_1 allows the extent of agent 2's underinvestment to be reduced, and by enough to offset any increase in underinvestment by agent 1. (We discuss the implications when states switch below.) A symmetric argument applies to the dynamics for high values of V_1 .

4.2. Long-run dynamics

To examine long-run convergence, we treat choices at date t as random variables and write $x_1(t)$ for the random value of net consumption of agent 1 at date t after history s^t etc. Define $\rho(t) := u'_2(x_2(t))/u'_1(x_1(t))$ to be the ratio of marginal utilities at date t ($\rho(t) = \sigma(t+1)$). In this subsection we focus on the long-run properties of $\rho(t)$.

With more than one state, convergence to constrained surplus maximization may not occur because there is a conflict between risk sharing and surplus maximization. To achieve surplus maximization in state s , the distribution of consumption may differ from that in $s' \neq s$ and therefore, an optimal contract must (dynamically) trade-off risk sharing against surplus maximization.

As already described, there is a (possibly trivial) interval of marginal utility ratios corresponding to maximum surplus in any state s . Let $[\underline{\rho}_s, \bar{\rho}_s]$ denote this interval in state s .¹⁴ By equation (4.1), the marginal utility ratio is unchanged from the previous period if (and only if) surplus is maximized today (i.e., $\rho(t) \in [\underline{\rho}_s, \bar{\rho}_s]$). Thus, a constant marginal utility ratio requires that $\Omega := \cap_{s \in \mathcal{S}} [\underline{\rho}_s, \bar{\rho}_s] \neq \emptyset$. The set Ω is non-empty when an FBA is sustainable, in which case the ratio is constant. If Ω is not only non-empty but a non-trivial interval, then there are multiple FBAs. Moreover, if an FBA is sustainable, then monotone convergence to an FBA occurs. If however, Ω is empty, or if CSM actions are not always efficient, an FBA is not sustainable and the marginal utility ratio may not converge to a single value. Nevertheless, under a weak regularity condition, it does converge to a unique long-run invariant distribution, independent of the initial conditions.

To describe the evolution of the marginal utility ratio, let $F_t^{(V_1(s_0))}: \mathbb{R}_+ \rightarrow [0, 1]$ denote the distribution function of $\rho(t)$ at date t given the initial value $V_1(s_0)$. This leads us to the following general convergence theorem.¹⁵

Theorem 1. (a) Suppose an FBA is sustainable. Then an optimal contract converges with probability one to an FBA: $\|a(t) - a^*(s_t)\| \rightarrow 0$ and the random sequence $\{\rho(t)\}$ is (weakly) monotone, with probability one. If there exist multiple FBAs, then the limit FBA depends upon $V_1(s_0)$.

(b) Suppose instead that an FBA is not sustainable. Then, provided $\pi_{ss} > 0$ for all s , $F_t^{(V_1(s_0))}$ converges weakly to a unique distribution independent of $V_1(s_0)$. Either (i) this distribution is degenerate, in which case dynamics are as in part (a), with stationary limit contract with CSM actions $\bar{a}(s)$ in each state, or otherwise (ii) this distribution is non-degenerate, and current surplus is not maximized in the long run: $\|a(t) - \bar{a}(s_t)\| \rightarrow 0$ with probability zero.

In part (a) of Theorem 1, there is convergence to an FBA. There is a (possibly trivial) interval of the ratio of marginal utilities given by the set Ω that are compatible with efficient actions

¹⁴ For $s \in \mathcal{S}_*$, $\underline{\rho}_s = \sigma_s(\hat{x}_1^s)$ and $\bar{\rho}_s = \sigma_s(\underline{x}_1^s)$ and for $s \in \mathcal{S}_*^c$, $\underline{\rho}_s = \bar{\rho}_s = \sigma_s(\hat{x}_1^s)$ (see Fig. 1).

¹⁵ We use $\|\cdot\|$ to denote the Euclidean norm.

and a constant marginal utility ratio. Convergence will be to the lower endpoint of Ω if the initial marginal utility ratio is below the interval; to the upper endpoint of Ω if initial marginal utility ratio is above the interval; and the sequence of marginal utility ratios will be constant if the initial marginal utility ratio belongs to Ω . The dynamics are similar in Part (b)(i), which considers the case where there is a marginal utility ratio consistent with CSM actions in each state. Convergence is to the CSM actions and to this (unique marginal utility) ratio. This case arises if there is a single state but CSM actions are not efficient.¹⁶ If there are multiple states and CSM actions in each state are inefficient, then this case is possible but not generic in the sense that a small perturbation of either ϕ_i^s or y^s in any state s will lead to the case of Part (b)(ii).

Part (b)(ii) of Theorem 1 provides a description of what happens when there is a conflict between surplus maximization and risk sharing. The optimal contract exhibits a second-best property. The marginal utility ratio $\rho(t)$ does not settle down to a single value, and whenever it differs across two dates $t-1$ and t , actions at date t will not be CSM.¹⁷ By contrast, in the risk-neutral case, we show that once the stationary phase is reached surplus is maximized in each state by varying the continuation utility to allow the constrained maximal surplus to be achieved (Theorem 2). For example, if the state changes from one in which agent 1 can claim most of output to one in which roles are reversed, sufficient surplus and future utility is reallocated to agent 2 to satisfy his self-enforcing constraint at the CSM actions for that state. However, in the risk-averse setting of part (b)(ii) of Theorem 1, risk-sharing considerations make such an immediate step change undesirable. It is better to hold agent 1's action at the later date inefficiently low, keeping agent 2's default payoff from rising too much, thereby relaxes the latter's self-enforcing constraint meaning that the share going to agent 2 does not rise to that consistent with the CSM actions.

To better understand this dynamic trade-off between surplus maximization and risk sharing suppose to the contrary that the ratio of marginal utilities differs across two dates $t-1$ and t , but actions at date t are CSM. Then a simple change in the contract at $t-1$ and t can produce a Pareto-improvement. Consider the case where $\rho(t-1) > \rho(t)$. Initially hold actions fixed at both dates and increase $x_1(t)$ by a small amount, but reduce $x_1(t-1)$ to leave $V_1(t-1)$ unchanged. If surplus were unchanged at t , this would improve risk sharing and lead to a Pareto-improvement because $V_2(t-1)$ would increase. However, because $x_2(t)$, and hence $V_2(t)$, have fallen, agent 2's self-enforcing constraint may be violated at the initial actions (and will be, if the CSM actions are below the first-best). In order not to violate agent 2's self-enforcing constraint, agent 1's action at date t can be reduced. Correspondingly, agent 2's action can be increased because $V_1(t)$ has risen. Critically, although this change may reduce surplus at date t , it does so only by a second-order amount since, by assumption, the original actions at date t were CSM.¹⁸ Consequently, a Pareto-improvement results, contradicting the supposed optimality of the original situation.

¹⁶ For the single state case irrespective of whether $s \in S_*$ or not, $\{\rho(t)\}$ monotone implies $\{z^s(t)\}$ is monotone increasing, converging to constrained maximal surplus, as indicated by the arrows in Fig. 1.

¹⁷ Formally, $\rho(t-1) \neq \rho(t)$ corresponds to $\sigma(t) \neq \sigma(t+1)$, and thus, from (4.1), $dz^{st}(V_1(t))/dV_1 \neq 0$. Hence, actions at date t are not CSM, as claimed.

¹⁸ The change in surplus would be second order when V_1 and V_2 are varied according to the Pareto frontier at t starting from maximum surplus; because the frontier's slope is $-\rho(t)$ at maximum surplus, the change we construct also only has a second-order effect. Also, note that, by construction, the self-enforcing constraints hold at t , and since $V_1(t-1)$ is unchanged and $V_2(t-1)$ is increased, they also hold at $t-1$.

4.3. Pure risk-sharing

We now compare our results to the standard limited commitment, two-agent, pure risk-sharing model of Thomas and Worrall (1988); Kocherlakota (1996); Ligon et al. (2002). To do this, for simplicity we consider a special case of our hold-up model with additive production, $y^s(a) = f_1^s(a_1) + f_2^s(a_2)$, and proportional defaults, $\phi_i^s(a) = \theta_{i1}^s f_1^s(a_1) + \theta_{i2}^s f_2^s(a_2)$ where $\theta_{ij}^s \geq 0$, $i, j = 1, 2$, and $\sum_{i=1}^2 \theta_{ij}^s = 1$, $j = 1, 2$. Our hold-up assumption requires $\theta_{ij}^s > 0$, $i, j = 1, 2$, $i \neq j$, all s . Holding technology and preferences fixed, consider the limit case where hold-up vanishes: $\theta_{ij}^s = 0$, $i, j = 1, 2$, $i \neq j$, all s . This corresponds to the pure-risk sharing model. In any optimal contract of this limit model actions are clearly efficient, as are actions in the breakdown, so only efficient levels play any role. Agent i 's "endowment" in state s is $f_i^s(a_i^*(s)) - a_i^*(s)$ and breakdown utility is $u(f_i^s(a_i^*(s)) - a_i^*(s))$.

We establish that the dynamics of the hold-up model converge to that of the risk-sharing model. In the latter, as is well known, dynamics are summarized in a simple updating rule for $\rho(t)$ (which fixes surplus division given surplus depends only on s). We characterize how the corresponding updating rule in the hold-up model converges to the risk-sharing one as hold-up disappears. One application of this is that it allows us to characterize general properties of the hold-up dynamics for cases where hold-up is low.

From Ligon et al. (2002), the updating rule in the pure risk-sharing case, which we write $\rho(t) = h^{RS}(\rho(t-1), s_t)$, has the property that there is a (possibly degenerate) interval $[\underline{\rho}_s^{RS}, \bar{\rho}_s^{RS}]$ for each s such that $h^{RS}(\rho(t-1), s) = \bar{\rho}_s^{RS}$ if $\rho(t-1) > \bar{\rho}_s^{RS}$; $h^{RS}(\rho(t-1), s) = \rho(t-1)$ if $\rho(t-1) \in [\underline{\rho}_s^{RS}, \bar{\rho}_s^{RS}]$ and $h^{RS}(\rho(t-1), s) = \underline{\rho}_s^{RS}$ if $\rho(t-1) < \underline{\rho}_s^{RS}$. Moreover, whenever optimal contracts that improve on autarky exist (if there is more than one distinct state, and δ is close enough to 1), each $[\underline{\rho}_s^{RS}, \bar{\rho}_s^{RS}]$ is non-degenerate (Proposition 2(iv) in Ligon et al., 2002).

Likewise, in the hold-up model we can also use $(\rho(t-1), s_t)$ as the state variable. (By $\rho(t-1) = \sigma(t)$, this is equivalent to $(\sigma(t), s_t)$.) Thus, the evolution of the contract can be represented by $\rho(t) = h(\rho(t-1), s_t)$, where $h: \mathbb{R}_+ \cup \{\infty\} \times \mathcal{S} \rightarrow \mathbb{R}_+$ (see Appendix A for details and characterization). The updating functions $h(\rho, s)$ converge to those of the pure risk-sharing model as the hold-up problem diminishes. Moreover, for $\rho(t-1)$ within the interior of the interval $[\underline{\rho}_{s_t}^{RS}, \bar{\rho}_{s_t}^{RS}]$, when hold-up is small enough, optimal actions at t are at the first-best levels and so $\rho(t) = \rho(t-1)$. An illustration of this convergence for two states is depicted in Fig. 2.¹⁹

Proposition 5. For each state $s \in \mathcal{S}$, (i) for all $\rho \in \mathbb{R}_+$, $h(\rho, s) \rightarrow h^{RS}(\rho, s)$ as $\theta_{ij} \rightarrow 0$, $i, j = 1, 2$, $i \neq j$, all s . (ii) If optimal contracts in the risk-sharing problem improve upon autarky, then for any η satisfying $(1/2)(\bar{\rho}_s^{RS} - \underline{\rho}_s^{RS}) > \eta > 0$, all s , there exists $\epsilon > 0$ such that for $\theta_{ij}^s < \epsilon$, $i, j = 1, 2$, $i \neq j$, all s , $h(\rho, s) = \rho$ for all $\rho \in [\underline{\rho}_s^{RS} + \eta, \bar{\rho}_s^{RS} - \eta]$.

One well-known feature of the pure risk sharing model is the "amnesia" property that once one of the agents is constrained, then the previous history is irrelevant to the future evolution of the optimal contract. This property no longer applies in our model of risk averse agents with actions. Suppose that agent 2's self-enforcing constraint binds at date t . In the risk-sharing problem, this

¹⁹ Fig. 2 is drawn with $h^{RS}(\rho, s)$ and $h(\rho, s)$ coinciding along the 45° line. This is for illustration only. Close to the limit as hold-up vanishes it is approximately true that the lines $h^{RS}(\rho, s)$ and $h(\rho, s)$ are coincident along the 45° line, except near the end-points of the segment of $h^{RS}(\rho, s)$ that lies on the 45° line. In general, we have no result on how the intersections of $h^{RS}(\rho, s)$ and $h(\rho, s)$ with the 45° line are related.




Fig. 2. Convergence to Pure Risk-Sharing.

fixes his continuation utility and there is a unique optimal way of delivering this continuation utility independently of past history and, in particular, independently of the previous ratio of marginal utilities. This can be seen in the flat sections of the functions $h^{RS}(\rho, s)$ in Fig. 2. In the hold-up problem, by contrast, agent 2's self-enforcing constraint can be relaxed by cutting agent 1's action. Although this change may reduce surplus, sacrificing surplus can be offset by improved risk sharing and the incentive to do this will vary with the lagged marginal utility ratio. The logic of trading off surplus to improve risk sharing is similar to the explanation given above for why the partial insurance case involves optimal actions that are not CSM, even in the long run. This result is illustrated in Fig. 2 by the fact that the functions $h(\rho, s)$ are upward sloping even away from the 45° line. Thus, even when an agent is constrained, past history affects the current actions and consumption and the future evolution of the optimal contract. The amnesia property fails.

5. Risk neutrality

For this section, we use Assumption A5 and suppose that both agents are risk neutral, in particular, that $u_i(x) = x$ and that $\underline{x}_i = -\infty$ for $i = 1, 2$. In this case, the non-negativity constraint on consumption (limited liability) plays a key role. We show that an optimal contract exhibits a two-stage property. It starts with a backloading phase in which one of the agents consumes all of the output. This agent never overinvests, while the other agent overinvests. The second phase is stationary and actions are CSM. Therefore, if $s \in \mathcal{S}_*$, actions are at the first-best for both agents. If $s \in \mathcal{S}_*^c$, both agents underinvest and have positive consumption. Depending on the initial division of surplus however, the optimal contract might start off in the stationary phase in which case the first backloading phase does not exist.

The lower bound for the deviation utility is strictly positive. Therefore, the Pareto-frontier is defined on $\Lambda^s := [\underline{V}_1^s, \bar{V}_1^s] \subset \mathbb{R}_{++}$. It can be shown that the frontier is strictly concave if at least one of the self-enforcing constraints is binding. If V_1 is in an interval where the efficient actions are sustainable (such values may not exist), then the frontier is linear with slope of -1 in this interval. In either case, CSM actions are unique.

Consider three (not necessarily disjoint) subsets of Λ^s : $A^s = \{V_1 \in \Lambda^s: c_1^o = 0\}$, $B^s = \{V_1 \in \Lambda^s: c_1^o > 0 \text{ and } c_2^o > 0\}$ and $C^s = \{V_1 \in \Lambda^s: c_2^o = 0\}$ where (c_1^o, c_2^o) represents an optimal value for consumption at V_1 . Note that $A^s \cup B^s \cup C^s = \Lambda^s$. Also note that A^s can be non-empty and C^s empty or vice-versa (examples of this type can be constructed). We know from our previous discussion that if agent 1 overinvests, this can only occur for $V_1 \in A^s$, and if agent 2 overinvests, this occurs for $V_1 \in C^s$. Also, since optimal actions are positive, output and aggregate consumption is positive, and consequently, it is not possible that both $\gamma_i > 0$ for the same V_1 . Equally, for $V_1 \in A^s$, $c_2 > 0$, and hence, the multiplier $\gamma_2 = 0$.²⁰ We also know from Proposition 1 that if $c_1 = 0$, and therefore, that agent 2 gets all the consumption, then agent 2 is unconstrained, and hence, $\mu_2 = 0$. Likewise, for $V_1 \in C^s$, $\gamma_1 = \mu_1 = 0$. Consumption for both agents is positive for $V_1 \in B^s$, so that $\gamma_1 = \gamma_2 = 0$.

Consider the subset A^s . Using $\gamma_2 = \mu_2 = 0$, we have from the first-order conditions (3.2a)–(3.2c) that:

$$\sigma_s^+(V_1) = 1 - \gamma_1 = \frac{\partial y^s(a_1, a_2)}{\partial a_1}, \tag{5.1a}$$

$$\sigma_s(V_1) = 1 - \gamma_1 - \mu_1 = \frac{\partial y^s(a_1, a_2)}{\partial a_1} - \frac{\partial z^s(a_1, a_2)}{\partial a_2} \left(\frac{dD_1^s}{da_2} \right)^{-1}. \tag{5.1b}$$

Hence, for $V_1 \in A^s$, $1 \geq \sigma_s^+(V_1) \geq \sigma_s(V_1)$. From equation (5.1a) it follows that if $\sigma_s^+(V_1) < 1$, then $\gamma_1 > 0$, and $\partial y^s(a_1, a_2)/\partial a_1 < 1$, so that agent 1 is overinvesting. From equation (5.1b) it follows that agent 2 doesn't overinvest and may underinvest. A similar set of conditions apply for $V_1 \in C^s$ and imply $1 \leq \sigma_s^+(V_1) \leq \sigma_s(V_1)$ so that agent 1 doesn't overinvest and if $\sigma_s^+(V_1) > 1$, then agent 2 overinvests. For $V_1 \in B^s$, the first-order conditions show that $\sigma_s^+(V_1) = 1$, so there is no overinvestment. As a measure of the extent of overinvestment let $\zeta_i^s := \max\{0, -\ln(\partial y^s(a_1, a_2)/\partial a_i)\}$ and $\zeta^s := \max\{\zeta_1^s, \zeta_2^s\}$. Hence, $\zeta^s > 0$ if there is overinvestment and is a measure of the distortion of the marginal product below the efficient level.²¹

We now state our two-phase characterization theorem. Here, for convenience, we also treat contracts as sequences of random variables, writing $a_i(t)$ rather than $a_i(s^t)$ etc.

Theorem 2. *In an optimal contract, there is a random time \hat{t} , $0 \leq \hat{t} < \infty$ with probability one, such that:*

Stationary phase ($t \geq \hat{t}$): *Optimal actions maximize the surplus $z^s(a_1, a_2)$ subject to the self-enforcing constraints, and hence, are CSM. The optimal actions depend only on the state s_t and are therefore independent of the initial conditions. There is no overinvestment: $a_i(t) \leq a_i^*(a_j(t), s_t)$ for $i, j = 1, 2, i \neq j$. For $s_t \in \mathcal{S}_*$, therefore, optimal actions and the corresponding surplus are first best: $a(t) = a^*(s_t)$ and $z^{s_t}(a(t)) = z^{s_t}(a^*(s_t))$. For $s_t \in \mathcal{S}_*^c$, the self-enforcing constraints bind for both agents, $c_i > 0$ for $i = 1, 2$, and there is underinvestment: $a_i(t) < a_i^*(a_j(t), s_t) \leq a_i^*(s_t)$ for $i, j = 1, 2, i \neq j$.*

Backloading phase ($t < \hat{t}$): *Overinvestment declines during the backloading phase: in particular, $\zeta(t)$ is weakly decreasing with $\zeta(\hat{t} - 1) = 0$. Backloading only applies to one agent, i , whose identity depends on the initial surplus split: this agent overinvests and has zero consumption*

²⁰ Since the multiplier is unique, the conclusion that $\gamma_2 = 0$ is valid even if V_1 also belongs to B^s or to C^s . The same argument can be made for the other subsets and multipliers.

²¹ In subset A^s , $\zeta^s = -\ln \sigma_s^+(V_1)$ and in subset C^s , $\zeta^s = \ln \sigma_s^+(V_1)$.

at each $t < \hat{t} - 1$. In the final period of backloading, at date $\hat{t} - 1$, there is no overinvestment: $a_i(\hat{t} - 1) \leq a_i^*(s_{\hat{t}-1})$, but $a_j(\hat{t} - 1) < a_j^*(s_{\hat{t}-1})$ for $j \neq i$. Moreover, if at any two dates t and $t' > t$ the same state s occurs, then underinvestment diminishes and surplus increases: $\partial y^s(a(t))/\partial a_j \geq \partial y^s(a(t'))/\partial a_j \geq 1$ for $\hat{t} - 1 \geq t' > t$ and $z^s(a(t')) \geq z^s(a(t))$ for $\hat{t} \geq t' > t$.

For a given value of agent 1's lifetime utility $V_1(s_0)$, there corresponds a value σ_0 . From Theorem 2, we can describe a typical path as follows. Suppose $\sigma_0 < 1$ (a symmetric argument applies if $\sigma_0 > 1$). Then one of two possible scenarios applies. Either $V_1(s_0) \in B^{s_0}$ or $V_1(s_0) \in A^{s_0}$. In the former case, $\hat{t} = 1$ and the contract moves to the stationary phase in each state at the next period. There is no overinvestment in this case. In the latter case, either $\zeta_1(0) = 0$ and $\hat{t} = 1$ as in the previous case, or $\zeta_1(0) > 0$ in which case $\hat{t} > 1$ and there is a backloading phase in which $c_1(t) = 0$ and agent 1 overinvests. Correspondingly, V_1 is sufficiently low that agent 1's self-enforcing constraint binds and agent 2 underinvests to avoid violating agent 1's self-enforcing constraint; by contrast V_2 is high enough that agent 2's self-enforcing constraint is slack.²² The basic intuition for the backloading result is familiar from other dynamic contracting models. The claim is that if agent 2 is unconstrained and underinvesting, then agent 1 has zero consumption at all previous dates, her payments are optimally backloaded into the future. The idea is that if agent 1 has positive consumption, then backloading her consumption allows her later constraints to be relaxed, which in turn means agent 2 can increase his future investment level without violating agent 1's constraint. Since agents are risk neutral they do not care about the timing of consumption flows (keeping the action plans fixed) if the expected discounted value is the same, but the backloading will permit future surplus to be increased, leading to a Pareto-improvement. Consumption is backloaded to the maximum extent possible, $c_1(t) = 0$ throughout the phase, allowing maximum surplus to be achieved as quickly as possible. Furthermore, by increasing $a_1(t)$ above $a_1^*(a_2(t), s)$, with the extra output being allocated to agent 2, additional backloading can be achieved, and for a small amount of overinvestment, the reduction in surplus is second-order.²³ Two novel results in the two-sided environment concerning the backloading phase are the over-investment by the agent whose utility is backloaded (although over-investment does not persist into the stationary phase), and the fact that despite the possibility that property rights might vary radically and persistently between states, only one of the agents will ever be subject to backloading.

That there is overinvestment in the backloading phase is perhaps surprising given the hold-up problem and given that the literature, mentioned in the Introduction, that considers the case where only one agent takes an action finds that there is never any overinvestment. Consider the one-sided case with only agent 1 taking an action. If agent 2 gets sufficient of the surplus to allow a_1 to be more than a_1^* without agent 2 wanting to deviate, then the optimal contract will be stationary with $a_1 = a_1^*$. The benefit from overinvestment is that it allows more backloading of agent 1's utility when $c_1 = 0$. In this case however there is no benefit, but an efficiency

²² This characterization applies so long as $\zeta_1(t) > 0$ and assuming agent 1's self-enforcing constraint binds with a positive multiplier. With more than one state, we cannot rule out the possibility that in some states deviation utilities are so low that the self-enforcing constraints may not bind even when $\sigma(t) < 1$. In this latter case, from (3.2c) and (3.2a), $a_2(t) = a_2^*(a_1(t), s)$ and $\sigma^+(t) = \sigma(t)$.

²³ The incentive to overinvest diminishes over time (as can be seen from (5.1a), $\sigma^+(t)$ approaches 1). Equally, if the same state recurs along the path, underinvestment diminishes as the self-enforcing constraint is relaxed. The combined effect is that surplus $z^s(a(t))$ increases, and reaches a maximum when $\sigma(t) = 1$.

cost, and backloading can only increase agent 2's incentive to renege in the future, potentially necessitating lower (inefficient) future actions by agent 1. Thus, there is no overinvestment.

6. Conclusion

In this paper, we have analyzed the dynamic properties of a relational contract between two agents both of whom undertake a costly investment or action that yields joint benefits. We have shown that optimal contracts exhibit different properties depending on whether agents are risk neutral or risk averse. In the risk-neutral case, actions may be either above or below the efficient level and actions and the division of the surplus converge monotonically to a stationary solution at which actions are constrained surplus maximizing (either both are first-best or both are below the first-best level). In the risk-averse case, we also establish a convergence result but convergence may or may not be monotonic depending on whether it is possible to sustain a first-best allocation or not. We have demonstrated that the optimal contract converges to the pure-risk sharing results of Kocherlakota (1996) as our hold-up problem vanishes.

In the risk-averse case there is an interesting trade-off between hold-up and risk-sharing. The hold-up problem creates an opportunity to relax the default constraint by lowering actions. This in turn allows more risk-sharing to be achieved without leading to default. It would be interesting to evaluate whether the gain in risk-sharing would ever be sufficient to offset the loss in surplus created by the original hold-up problem. This is a difficult question because without additional structure to the model little can be said about the long run distribution of the optimal contract.

Acknowledgments

Funding: This work was supported by the Economic and Social Research Council [Research Grant: ES/L009633/1]. We thank the referees and editor for helping to improve the paper. In particular, we'd like to thank, Pierre Dubois, Francisco Gonzalez, Bruno Jullien and John Moore for very valuable comments. All errors are ours.

Appendix A

A.1. Statements of lemmas for Section 2

Lemma 1. Under Assumption 2, for $i, j = 1, 2, i \neq j$ and for each $s \in \mathcal{S}$, the conditionally efficient action, $a_i^*(a_j, s)$, is single-valued, weakly increasing and continuous in a_j .

Lemma 2. Under Assumption 3, for $i, j = 1, 2, i \neq j$ and for each $s \in \mathcal{S}$, the Nash best-response, $a_i^N(a_j, s)$, is single-valued, weakly increasing and continuous in a_j . Moreover, $0 < a_i^N(a_j, s) < a_i^*(a_j, s)$ for all a_j .

Lemma 3. Under Assumptions 1 and 3, for $i, j = 1, 2, i \neq j$ and for each $s \in \mathcal{S}$, the deviation utility, $D_i^s(a_j)$, is bounded below and is a continuous, increasing, strictly concave, and differentiable function of a_j .

A.2. Statements of lemmas and technical details for Section 3

Lemma 4. Under Assumptions 1–3, the set of lifetime utilities \mathcal{V}_{s_0} that correspond to dynamic relational contracts is compact for each $s_0 \in \mathcal{S}$. Hence, optimal contracts exist.

It is convenient in analyzing the recursive problem to change variables and use the deviation utilities of the two agents instead of actions. Let $d_j := D_i^s(a_j)$; by Lemma 3 $dD_i^s(a_j)/da_j > 0$, and we let $g_i^s(d_j) := (D_i^s)^{-1}(d_j)$. Abusing notation, surplus is $z^s(d_1, d_2) := z^s(g_2^s(d_1), g_1^s(d_2))$, with output $y^s(d_1, d_2)$ defined similarly. Given the properties of $D_i^s(a_j)$ (Lemma 3), the functions $g_j^s(d_i)$ are continuously differentiable, strictly increasing and strictly convex. Let $\mathcal{D}(s) := \{(d_1, d_2) = (D_2^s(a_1), D_1^s(a_2)) \mid (a_1, a_2) \in \mathbb{R}_+^2\}$. The contract $\{d(s^t), x(s^t)\}_{t=0}^\infty$ is feasible if $\sum_i x_i(s^t) \leq z^{s^t}(d(s^t))$ (total consumption does not exceed output) for every history s^t , and for actions and consumption to be non-negative, it must also satisfy $d(s^t) \in \mathcal{D}(s^t)$ for every history s^t and $x_i(s^t) + g_j^{s^t}(d_i(s^t)) \geq 0$ for $i, j = 1, 2, i \neq j$ and every history s^t . We define $d_i^*(s) = D_j^s(a_i^*(s)), i \neq j, d_i^*(d_j, s) = D_j^s(a_i^*(g_i^s(d_j), s))$, etc.

Problem [P1] can be reformulated with $d \in \mathcal{D}(s)$ replacing $a \geq 0$ as a choice variable, the RHS of (3.1b) and (3.1c) being d_2 and d_1 respectively, a_i in (3.1f) being $g_j^s(d_i)$ and the RHS of (3.1g) being $z^s(d_1, d_2)$, with solution denoted by $(d^s(V_1), x^s(V_1))$.²⁴

With the change in variables, the first-order condition (3.2c) becomes

$$\frac{\mu_j}{1 + \mu_2} = \frac{\partial z^s}{\partial d_i} \left(u'_2(\cdot) + \frac{\gamma_2}{1 + \mu_2} \right) + g_j^{s'}(d_i) \frac{\gamma_i}{1 + \mu_2} \quad i, j = 1, 2 \text{ and } i \neq j. \tag{A.1}$$

To establish concavity of $V_2^f(\cdot)$ we give two *alternative* assumptions.

Assumption A4. (a) $z^s(d) : \mathcal{D}(s) \rightarrow \mathbb{R}$ is strictly concave in d and (b) any solution to [P1] has $x_i > 0$ for $i = 1, 2$ and for each s^t .

Assumption A5. The function $z^s(d) + g_j^s(d_i) : \mathcal{D}(s) \rightarrow \mathbb{R}_+$ is concave in d for each $i, j = 1, 2, j \neq i$.

Under either Assumption A4 or Assumption A5, the Pareto-frontier is concave on $[V_1^s, \bar{V}_1^s]$ (Proposition 2(i)). Under Assumption A4, it is easily checked that the constraint set is also convex.²⁵

Although Assumptions A4 and A5 are not directly on primitives of the model (because they are specified in terms of the deviation utility and an endogenous variable for Assumption A4), it is easily checked that there are natural parameterizations of the model where these assumptions are satisfied. For example, Assumption A4 is satisfied provided that agents are not too risk averse. For example, consider the case where preferences exhibit constant absolute risk aversion with coefficient $\alpha > 0$, the same for both agents, and the production function is separable and given by $y^s(a_1, a_2) = (\beta)^{-1}((a_1)^\beta + (a_2)^\beta)$ where $\beta \in (0, 1)$. Furthermore, suppose each agent can

²⁴ The linear independence constraint qualification holds unless the constraints (3.1f) are inactive and $u'_2(\partial z^s / \partial a_1)(dg_2^s(d_1)/dd_1) = 0$. This constraint qualification can fail, but it only fails at $V_1 = \bar{V}_1^s$ where the slope of the Pareto-frontier is infinite (examples where the constraint qualification fails at this point can be constructed). Thus, apart from $V_1 = \bar{V}_1^s$, the linear independence constraint qualification holds and the Lagrangian multipliers in the first-order conditions (reported in sub-section 3.4) exist and are unique. We can also ignore points $V_1 = \bar{V}_1^s$ without loss of generality: if $V_1(s_0) < \bar{V}_1^{s_0}$, then we will show that $V_1 \neq \bar{V}_1^s$ for any state s ; if $V_1(s_0) = \bar{V}_1^{s_0}$, then it will be possible to reformulate the problem maximizing the utility of agent 1 for a given V_2 for agent 2 and the relevant constraint qualification will be satisfied.

²⁵ It can also be checked that if [P1] is written with c and d as choice variables, then a sufficient condition for convexity of the constraint set is that $y^s(d)$ is concave in d . This condition is more stringent than concavity of $z^s(d)$ and will fail in a number of natural cases.

expropriate a proportion θ of output in the case of default. Then a sufficient condition for the assumption to be satisfied for $\theta \in (1/e, 1/2]$ is if $\alpha < -e\theta(1 - \theta)^{-1} \log \theta$, and for $\theta \in (0, 1/e]$ if $\alpha < (1 - \theta)^{-1}$. Equally, suppose that agents are risk neutral with $u_i(x) = x$, production is additive and the breakdown consumption in each state is $\phi_i(a) = \theta_{i1} f_1(a_1) + \theta_{i2} f_2(a_2)$, where for notational simplicity the dependence of θ, f etc. on s is suppressed. With this specification for $\phi_i(a)$, $D''_j/D'_j = f''_j/f'_j$, and it can be checked that Assumption A5 is satisfied.

Lemma 5. Under Assumptions 1–3 and under either Assumption A4 or Assumption A5, $d^s_i(V_1)$ is a continuous function of V_1 for each $s \in \mathcal{S}$ and $i = 1, 2$.

Lemma 6. Under Assumptions 1–3, and for $i, j = 1, 2, i \neq j$, for any history s^t , (i) if $V_i(s^t) > d_j(s^t)$, then $a_j(s^t) \geq a^*_j(a_i(s^t), s_t)$; (ii) if $c_i(s^t) > 0$, then $a_i(s^t) \leq a^*_i(a_j(s^t), s_t)$.

A.3. Statement of lemmas and proofs of theorem for Section 4

For all lemmas and proofs in this subsection, we maintain Assumption A4. Additionally it is assumed that agents are risk averse, that is, u_i is strictly concave for $i = 1, 2$.

Lemma 7. For each $s \in \mathcal{S}$, a solution to [P1] has the property that $z^s(a_1, a_2)$ is maximized over $a \in \mathbb{R}^2_+$ subject to $V_1 \geq D^s_1(a_2)$ and $V_2^s(V_1) \geq D^s_2(a_1)$.

Lemma 8. For each $s \in \mathcal{S}$, the surplus function $z^s(V_1)$ is continuous, concave and differentiable in V_1 .

Lemma 9. For each $s \in \mathcal{S}$, (i) $dz^s(V_1)/dV_1 > 0$ (< 0) implies $\mu^s_1(V_1) > 0$ ($\mu^s_2(V_1) > 0$); (ii) there are two critical values $\bar{\chi}^s_1 \in (\underline{V}^s_1, \bar{V}^s_1]$ and $\underline{\chi}^s_1 \in [\underline{V}^s_1, \bar{V}^s_1)$, such that $d^s_2(V_1) = V_1$ for all $V_1 \leq \bar{\chi}^s_1$ and $d^s_1(V_1) = V_2^s(V_1)$ for all $V_1 \geq \underline{\chi}^s_1$. Moreover, $\mu^s_1(V_1) = 0$ for $\bar{V}^s_1 > V_1 \geq \bar{\chi}^s_1$ and $\mu^s_2(V_1) = 0$ for $\underline{V}^s_1 < V_1 \leq \underline{\chi}^s_1$ (if such V_1 exist). If the efficient actions can be sustained in state s , then $\bar{\chi}^s_1 \leq \underline{\chi}^s_1$. Otherwise, $\bar{\chi}^s_1 > \underline{\chi}^s_1$, and surplus is maximized for a unique value of $V_1 \in (\underline{\chi}^s_1, \bar{\chi}^s_1)$ at which both constraints bind.

Proof of Theorem 1:

Before proving the theorem, we prove two lemmas.

Since $\sigma^+_s(V_1)$ depends only on the current slope σ and the current state s (recall V_1 and σ are uniquely related for a given state) the evolution of the contract can be represented as a stochastic recursion, i.e., $\sigma(t + 1) = \sigma^+_s(\sigma^-(\sigma))$, which we write as $\sigma(t + 1) = h(\sigma(t), s_t)$, and where $h: \mathbb{R}_+ \cup \{\infty\} \times \mathcal{S} \rightarrow \mathbb{R}_+$; $\sigma(0) = \sigma_0$ is the given initial value, corresponding to the initial state s_0 and agent 1's lifetime utility $V_1(s_0)$. (This is the same function as h defined in the text, given that $\rho(t) = \sigma(t + 1)$.)

Lemma 10. (i) The function $h(\sigma, s)$ is continuous and strictly increasing in σ ; (ii) for each state s , there is a single, possibly degenerate, interval of fixed points $[\underline{\sigma}^*_s, \bar{\sigma}^*_s]$, $\underline{\sigma}^*_s > 0$, such that $h(\sigma, s) = \sigma$ for any $\sigma \in [\underline{\sigma}^*_s, \bar{\sigma}^*_s]$; (iii) $h(\sigma, s) < \sigma$ for $\sigma > \bar{\sigma}^*_s$ and $h(\sigma, s) > \sigma$ for $\sigma < \underline{\sigma}^*_s$.

Proof. Let $x_i(V_1) = c_i(V_1) - g_j(d_i(V_1))$ be the net consumption of agent i (dropping the state superscript) and $\rho(V_1) := u'_2(x_2(V_1))/u'_1(x_1(V_1))$. Then, $h(\sigma, s) = \rho(\sigma^{-1}_s(\sigma))$.

We first prove part (i). From the concavity properties of [P1] under Assumption A4, the choice variables $x_i(V_1)$ are continuous, and hence, $\rho(V_1)$ is continuous in V_1 . From Proposition 2(i), the Pareto-frontier is continuously differentiable and hence, so too is its inverse. Thus, $h(\sigma, s)$ is continuous in σ .

Next, we turn to the monotonicity of $h(\sigma, s)$. First, we show that $\rho(V_1)$ is strictly increasing. Suppose, to the contrary, that $\rho(V_1) \leq \rho(\tilde{V}_1)$ for some $V_1 > \tilde{V}_1$. It follows from $\rho(V_1) = -V_2^{r'}(V_1)$ and the concavity of the frontier $V_2^r(V_1)$ that $V_1^r(V_1) \leq V_1^r(\tilde{V}_1)$ for all $r \in \mathcal{S}$. Also, since $\tilde{V}_1 < V_1$, we have $u_1(x_1(\tilde{V}_1)) + \delta \sum_{r \in \mathcal{S}} \pi_{sr} V_1^r(\tilde{V}_1) < u_1(x_1(V_1)) + \delta \sum_{r \in \mathcal{S}} \pi_{sr} V_1^r(V_1)$. Hence, $x_1(\tilde{V}_1) < x_1(V_1)$. Likewise, since the frontier is downward sloping, $V_2(\tilde{V}_1) > V_2(V_1)$ and $V_2^r(V_1^r(\tilde{V}_1)) \leq V_2^r(V_1^r(V_1))$, and therefore, that $x_2(\tilde{V}_1) > x_2(V_1)$. But then $u_2'(x_2(\tilde{V}_1))/u_1'(x_1(\tilde{V}_1)) < u_2'(x_2(V_1))/u_1'(x_1(V_1))$ or $\rho(\tilde{V}_1) < \rho(V_1)$, which is a contradiction. Thus, we can conclude that $\rho(V_1)$ is strictly increasing in V_1 . Since the frontier $V_2^s(V_1)$ is strictly decreasing in V_1 and $\sigma = -V_2^{s'}(V_1)$, the result is proved.

To establish parts (ii) and (iii), from Lemma 9, for $s \in \mathcal{S}_*$, surplus is at the first-best level for $V_1 \in [\bar{\chi}_1^s, \underline{\chi}_1^s]$. Correspondingly, there is an interval of Pareto frontier slopes $[\underline{\sigma}_s^*, \bar{\sigma}_s^*] := [-V_2^{s'}(\bar{\chi}_1^s), -V_2^{s'}(\underline{\chi}_1^s)]$. For $s \in \mathcal{S}_*^c$, the corresponding interval is degenerate at a single point $[\underline{\sigma}_s^*, \bar{\sigma}_s^*] := [-V_2^{s'}(\hat{\chi}_1^s)]$ where $\hat{\chi}_1^s = \arg \max_{V_1} z^s(V_1)$. It follows from part (b) of Assumption A4 and equation (3.2b) (given $\gamma_i = v_i^r = 0$ for $i = 1, 2$) that $\sigma_s^+(V_1)$ is positive and finite. Thus, $0 < \underline{\sigma}_s^* \leq \bar{\sigma}_s^* < \infty$. Equation (4.1) therefore implies the following: If $\sigma_s(V_1) \in [\underline{\sigma}_s^*, \bar{\sigma}_s^*]$, then $\sigma_s^+(V_1) = \sigma_s(V_1)$. If $\sigma_s(V_1) > \bar{\sigma}_s^*$, then $dz^s(V_1)/dV_1 < 0$ (given the concavity of $z^s(V_1)$ by Lemma 8) and hence, $\sigma_s^+(V_1) < \sigma_s(V_1)$. Likewise, if $\sigma_s(V_1) < \underline{\sigma}_s^*$, then $dz^s(V_1)/dV_1 > 0$ and hence, $\sigma_s^+(V_1) > \sigma_s(V_1)$. \square

Let \bar{s} be a state such that $\underline{\sigma}_{\bar{s}}^* \geq \bar{\sigma}_{\bar{s}}^*$ and \underline{s} a state such that $\bar{\sigma}_{\underline{s}}^* \leq \underline{\sigma}_{\underline{s}}^*$ for all $s \in \mathcal{S}$.

Lemma 11. *An FBA is sustainable if and only if $\underline{\sigma}_{\bar{s}}^* \leq \bar{\sigma}_{\underline{s}}^*$ and $\mathcal{S}_* = \mathcal{S}$.*

Proof. The “if” implication follows because there would exist an initial value $\sigma_0 \in [\underline{\sigma}_{\bar{s}}^*, \bar{\sigma}_{\bar{s}}^*]$ such that $\sigma_0 \in [\underline{\sigma}_s^*, \bar{\sigma}_s^*]$ for each state s . It therefore follows that starting from σ_0 , $\sigma(t)$, and hence, the ratio of marginal utilities, is kept constant at σ_0 and since surplus is maximized for $\sigma(t) \in [\underline{\sigma}_{s_t}^*, \bar{\sigma}_{s_t}^*]$, actions are CSM and thus first-best by $\mathcal{S}_* = \mathcal{S}$ in each state. “Only if” follows because by Lemma 10 even if first-best actions are sustainable in every state, $\underline{\sigma}_{\bar{s}}^* > \bar{\sigma}_{\underline{s}}^*$ would imply that whenever $s_t = \bar{s}$ and $s_\tau = \underline{s}$ (such t, τ exist with probability one given irreducibility), then either (a) $\sigma(t) \in [\underline{\sigma}_{s_t}^*, \bar{\sigma}_{s_t}^*]$ and $\sigma(\tau) \in [\underline{\sigma}_{s_\tau}^*, \bar{\sigma}_{s_\tau}^*]$ in which case $\sigma(t) > \sigma(\tau)$, and the risk-sharing condition fails, or (b) either or both $\sigma(t) \notin [\underline{\sigma}_{s_t}^*, \bar{\sigma}_{s_t}^*]$ and $\sigma(\tau) \notin [\underline{\sigma}_{s_\tau}^*, \bar{\sigma}_{s_\tau}^*]$, in which case surplus is not maximized at least one of the dates. \square

Proof of Theorem 1. Recalling that $\rho(t) = \sigma(t+1)$, an interval $[\underline{\sigma}_s^*, \bar{\sigma}_s^*]$ corresponds to an interval of marginal utility ratios in state s and convergence of $\sigma(t)$ is equivalent to convergence of $\rho(t)$. Part (a) of the Theorem therefore follows straightforwardly from Lemmas 10 and 11. From Lemma 11 $\underline{\sigma}_{\bar{s}}^* \leq \bar{\sigma}_{\underline{s}}^*$. Convergence is to $\underline{\sigma}_{\bar{s}}^*$ if $\sigma_0 < \underline{\sigma}_{\bar{s}}^*$, since $\sigma(t) = h(\sigma(t-1), s_{t-1}) \geq \sigma(t-1)$ by Lemma 10 (iii) and so $\{\sigma(t)\}$ is a non-decreasing sequence; it is bounded above by $\underline{\sigma}_{\bar{s}}^*$ given h continuous and increasing in σ and that $h(\sigma, s) \leq \sigma$, all s , for $\sigma \geq \underline{\sigma}_{\bar{s}}^*$, by Lemma 10 (ii) and (iii); with probability one $\sigma(t)$ converges to $\underline{\sigma}_{\bar{s}}^*$ given that $h(\sigma, \bar{s}) > \sigma$ for $\sigma < \underline{\sigma}_{\bar{s}}^*$ and the irreducibility of $\{\pi_{sr}\}$ and finiteness of states implies that state \bar{s} is recurrent. Likewise convergence (monotonic) is to $\bar{\sigma}_{\underline{s}}^*$ if $\sigma_0 > \bar{\sigma}_{\underline{s}}^*$, and $\sigma(t)$ is constant at σ_0 if $\sigma_0 \in [\underline{\sigma}_{\bar{s}}^*, \bar{\sigma}_{\underline{s}}^*]$. If there

exist multiple FBAs then $\underline{\sigma}_s^* < \bar{\sigma}_s^*$, and the limit depends on σ_0 and hence on $V_1(s_0)$. Since $\lim_{t \rightarrow \infty} \sigma(t) \in [\underline{\sigma}_s^*, \bar{\sigma}_s^*]$ for all s , $\bar{\rho}(t)$ converges and by continuity the limit actions are $a^*(s_t)$, and $x^*(s_t)$ is such that $u_2^*(x_2^*(s))/u_1^*(x_1^*(s)) = \lim_{t \rightarrow \infty} \sigma(t)$, all s .

For (b), if an FBA is not sustainable, then by Lemma 11 either $\underline{\sigma}_s^* = \bar{\sigma}_s^*$ and the CSM actions are below first-best levels in at least one state, or $\underline{\sigma}_s^* > \bar{\sigma}_s^*$. In the former case by Lemma 10, following the argument in part (a), $\{\sigma(t)\}$ is monotonic and converges with probability one to $\underline{\sigma}_s^* = \bar{\sigma}_s^* \in [\underline{\sigma}_s^*, \bar{\sigma}_s^*]$ for all s , implying that limiting actions are CSM in each state, establishing case (b)(i). Otherwise, part (b)(ii) obtains; $\sigma(t) \in [\min_s \{h(0, s)\}, \max_s \{h(\infty, s)\}]$ for $t \geq 1$. Irreducibility and finiteness of $\{\pi_{s_t}\}$ implies \underline{s} and \bar{s} are recurrent, and $\underline{\sigma}_s^* > \bar{\sigma}_s^*$ implies that for any $\sigma \geq 0$, either $h(\sigma, \bar{s}) > \sigma$ or $h(\sigma, \underline{s}) < \sigma$ (or both). Thus, given h is continuous in σ , weak convergence to a degenerate distribution is impossible. Next consider the sequence of r.v.s $\{a(t) - \bar{a}(s_t)\}$. Assume w.l.o.g. that state \underline{s} is uniquely defined. Consider an infinite history $\{s_0, s_1, \dots\}$ in which each state occurs infinitely often, which implies from the properties of h established in Lemma 10 that there exists t' such that $\sigma(t) \geq \bar{\sigma}_s^*$ for $t \geq t'$; note that the set of such histories has probability one. Suppose that $a(t) - \bar{a}(s_t) \rightarrow 0$, so that along the subsequence $\{s_{t_1}, s_{t_2}, s_{t_3}, \dots\}$ where t_i is the i th time \underline{s} occurs, $a(t_i) \rightarrow \bar{a}(\underline{s})$ as $i \rightarrow \infty$. Consider a $t \geq t'$ such that $s_t = \bar{s}$. Then $\sigma(t) \geq h(\bar{\sigma}_s^*, \bar{s}) > \bar{\sigma}_s^*$ by h increasing in σ and $h(\sigma, \bar{s}) > \sigma$ for $\sigma < \bar{\sigma}_s^*$. If t_i is the next time \underline{s} occurs, $\sigma(t_i) \geq \min\{h(\bar{\sigma}_s^*, \bar{s}), \min_{s \neq \underline{s}} \bar{\sigma}_s^*\} > \bar{\sigma}_s^*$. This implies that $V_1(t_i)$ is bounded above $\arg \max_{V_1} z^s(V_1)$, i.e., above $\sigma_s^{-1}(\bar{\sigma}_s^*)$, so $a(t_i)$ is bounded away from $\bar{a}(\underline{s})$. Since $s_t = \bar{s}$ infinitely often, this contradicts $a(t_i) \rightarrow \bar{a}(\underline{s})$. Next, fix any $\sigma_c \in (\bar{\sigma}_s^*, \underline{\sigma}_s^*)$; clearly $h(\sigma_c, \bar{s}) > \sigma_c$ and $h(\sigma_c, \underline{s}) < \sigma_c$. Using $\pi_{s_s} > 0$ all s , there exist $t \geq 1$ such that

$$\begin{aligned} \varepsilon_1 &:= \mathbb{P}(\sigma(t) < \sigma_c \mid \sigma_0 = (\max_s \{h(\infty, s)\})) > 0 \\ \varepsilon_2 &:= \mathbb{P}(\sigma(t) > \sigma_c \mid \sigma_0 = (\min_s \{h(0, s)\})) > 0, \end{aligned}$$

since for ε_1 (respectively ε_2) consider a sufficient number of consecutive occurrences of \underline{s} (respectively \bar{s}). This implies the “splitting condition” of Bhattacharya and Majumdar (2007, Chapter 3.5, p. 250) for the i.i.d. case, and the condition in Foss et al. (2018, Corollary 1) in the general Markov case. Thus, there is a unique stationary distribution \tilde{F} such that $F_t^{(V(s_0))}$ converges weakly to \tilde{F} , as $t \rightarrow \infty$, for any initial condition. \square

A.4. Statement of lemmas and proof of theorem for Section 5

For this subsection we maintain Assumption A5 but additionally assume that agents are risk-neutral with $u_i(x_i) = x_i$.

Lemma 12. For each $s \in \mathcal{S}$, the Pareto-frontier $V_2^s(\cdot)$ is strictly concave on $[\underline{V}_1^s, \underline{V}_1^{s*})$ where $\underline{V}_1^{s*} := \inf\{V_1 : V_2^{s'}(V_1) = -1\}$, and on $(\bar{V}_1^{s*}, \bar{V}_1^s]$ where $\bar{V}_1^{s*} := \sup\{V_1 : V_2^{s'}(V_1) = -1\}$. If first-best actions are not sustainable in state s , i.e., for $s \in \mathcal{S}_*^c$, then $V_2^s(\cdot)$ is strictly concave on $[\underline{V}_1^s, \bar{V}_1^s]$.

Lemma 13. With probability one, there is a random time $\hat{t} < \infty$ such that $\zeta(t)$ converges monotonically to 0 with $\zeta(t) = 0$ for all $t \geq \hat{t} - 1$.

Lemma 14. For each $s \in \mathcal{S}$, the surplus function $z^s(V_1)$ is a continuous and single peaked function of V_1 . That is, for any $V_1^{(1)} < V_1^{(2)} < V_1^{(3)}$, it is not possible that $z^s(V_1^{(1)}) > z^s(V_1^{(3)}) > z^s(V_1^{(2)})$. Moreover, $z^s(V_1)$ is maximal when $\sigma_s(V_1) = 1$.

Proof of Theorem 2. Stationary phase: Define \hat{t} as the earliest date at which $\sigma(t) = 1$. By Lemma 13, $\hat{t} < \infty$ with probability 1. If $\sigma = 1$ and $V_1 \in A$ (so that $c_1 = 0$), then it follows from the first-order conditions that $\gamma_1 = \gamma_2 = \mu_1 = \mu_2 = 0$, and hence that $a(t) = a^*(s_t)$ or equivalently $d(t) = d^*(s_t)$ (we suppress t below for notational simplicity). A similar argument applies for $\sigma = 1$ and $V_1 \in C$. If $\sigma = 1$ and $V_1 \in B$ (so $c_1 > 0$ and $c_2 > 0$), then $\mu_1 = \mu_2$ from (3.2a). Thus, either $\mu_i = 0, i = 1, 2$, in which case again $d = d^*$ and $z^s(d)$ is maximal, or $\mu_i > 0, i = 1, 2$, so both self-enforcing constraints bind. In the latter case, $(\partial z^s / \partial d_2) / (\partial z^s / \partial d_1) = 1 = -V_2^s(V_1)$, from (3.2a) and (A.1). From the concavity of $V_2^s(V_1)$ and $z^s(d)$, this implies that $z^s(d)$ is maximized by choice of $d \in \mathcal{D}(s)$ and $V_1 \in [\underline{V}_1^s, \bar{V}_1^s]$ subject to $d_2 \leq V_1, d_1 \leq V_2(V_1)$. Since these self-enforcing constraints must hold for any dynamic relational contract, it follows that at $\sigma = 1, z^s(d)$ is maximal across all dynamic relational contracts whether the self-enforcing constraints bind or not and optimal actions are CSM. Also in the case where $\mu_1 = \mu_2 > 0$, it follows from (A.1) that $d_i < d_i^*(d_j, s) \leq d_i^*(s), j \neq i$ (the last inequality follows because $d_i^*(d_j, s)$ is non-decreasing in d_j). Since a_i and d_i are positively monotonically related through the function g_j^s and $a = a^*$ if and only if $d = d^*$, the statement in the Theorem follows.

Backloading phase: Suppose $V_1(s_0)$ is such that $\sigma_0 < 1$ (a symmetric argument applies if $\sigma_0 > 1$). Then $\sigma(t) \leq \sigma^+(t) \equiv \sigma(t+1) \leq 1$ and $\gamma_2(t) = 0$.

We first establish the last part of the theorem. Consider $t = \hat{t} - 1$, so that $\sigma^+(t) \equiv \sigma(t+1) = 1$. Then $\sigma^+(t) - \sigma(t) > 0$ and from (3.2a), $\mu_1(t) > 0$. (3.2b) implies that $\gamma_1(t) = 0$. So from (A.1), $\partial z^{s_t} / \partial d_1 \geq 0$, and thus, $d_1(t) \leq d_1^*(d_2(t), s_t)$. Likewise in (A.1), $d_2(t) < d_2^*(d_1(t), s_t)$. Together with $d_1(t) \leq d_1^*(d_2(t), s_t)$ this implies $d_1(t) \leq d_1^*(s_t)$ and $d_2(t) < d_2^*(s_t)$; equivalently $a_1(t) \leq a_1^*(s_t)$ and $a_2(t) < a_2^*(s_t)$.

Next, suppose $\sigma^+(t) < 1$, so that $t < \hat{t} - 1$ and from (5.1a) $\gamma_1 > 0$ (so $c_1 = 0$ and $V_1(t) \in A^{s_t}$). Equation (5.1a) implies that $\partial y^s(a_1, a_2) / \partial a_1 < 1$, so that $a_1(t) > a_1^*(a_2(t), s_t)$. Again using (A.1), $\partial z^s(d_1, d_2) / \partial d_2 \geq 0$, so that $d_2(t) \leq d_2^*(d_1(t), s_t)$ and hence $a_2(t) \leq a_2^*(a_1(t), s_t)$.

To establish the monotonicity of the marginal conditions, consider dates t and t' with $\hat{t} \geq t' > t$ such that the same state s occurs at date t and t' . If $\sigma^+(t') < 1$, then it follows from the monotonicity of the sequence established in Lemma 13 that $\sigma^+(t) \leq \sigma^+(t') < 1$. Hence, $V_1(t) \in A^{s_t}$ and $V_1(t') \in A^{s_{t'}}$. It follows directly from (5.1a) that $1 > \partial y^s(a(t')) / \partial a_1 \geq \partial y^s(a(t)) / \partial a_1$. Now suppose, contrary to the assertion that $\partial y^s(a(t')) / \partial a_2 > \partial y^s(a(t)) / \partial a_2$ or equivalently $\partial z^s(a(t')) / \partial a_2 > \partial z^s(a(t)) / \partial a_2$. From (5.1a) and (5.1b) $\partial z^s(a(t)) / \partial a_2 (dD_1^s / da_2)^{-1} \equiv \partial z^s(d(t)) / \partial d_2 = \sigma^+(t) - \sigma(t) \geq 0$, and hence $\partial z^s(a(t)) / \partial a_2 \geq 0$ and $\partial z^s(a(t')) / \partial a_2 > 0$. Strict concavity of $z^s(a)$ requires that $\sum_{i=1}^2 ((\partial z^s(a(t')) / \partial a_i) - (\partial z^s(a(t)) / \partial a_i))(a_i(t') - a_i(t)) < 0$ for $a(t) \neq a(t')$. Then, since $\partial^2 z^s / \partial a_1 \partial a_2 \geq 0$, it follows that $a_1(t) \geq a_1(t')$ and $a_2(t) > a_2(t')$. This however provides a contradiction. To see this, consider that $\sigma(t') \geq \sigma(t)$ and $\sigma(t) < 1$ imply, from Lemma 12, that $V_1(t') \geq V_1(t)$. Equally, because $\partial z^s(a(t')) / \partial a_2 > 0, \mu_1(t') > 0$, and hence, $d_2(t') = V_1(t') \geq V_1(t) \geq d_2(t)$. This implies $a_2(t') \geq a_2(t)$, a contradiction. A similar argument applies if $\sigma^+(t) < \sigma^+(t') = 1$ except that in this case we have $\gamma_1(t') = \gamma_2(t') = 0$ and thus, from (A.1), $\partial z^s(a(t')) / \partial a_1 \geq 0 > \partial z^s(a(t)) / \partial a_1$ (the second inequality follows from the earlier argument because $\sigma^+(t) < 1$). Finally, Lemma 14 shows that $z^s(V_1)$ is continuous and single-peaked and has a maximum when $\sigma_s(V_1) = 1$. From above, $V_1(t') \geq V_1(t)$ when $\sigma(t) < 1$. Hence, we conclude that $z^s(a(t)) \leq z^s(a(t'))$. \square

Appendix B. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jet.2018.02.004>.

References

- Ábrahám, Árpád, Lacsó, Sarolta, 2018. Efficient risk sharing with limited commitment and storage. *Rev. Econ. Stud.* <https://doi.org/10.1093/restud/rdx061>. Forthcoming.
- Acemoğlu, Daron, Golosov, Mikhail, Tsyvinski, Aleh, 2011. Power fluctuations and political economy. *J. Econ. Theory* 146 (3), 1009–1041. <https://doi.org/10.1016/j.jet.2010.11.002>.
- Albuquerque, Rui, Hopenhayn, Hugo A., 2004. Optimal lending contracts and firm dynamics. *Rev. Econ. Stud.* 71 (2), 285–315. <https://doi.org/10.1111/0034-6527.00285>.
- Bhattacharya, Rabi, Majumdar, Mukul, 2007. *Random Dynamical Systems: Theory and Applications*. Cambridge University Press, Cambridge. ISBN 978-0521532723.
- Broer, Tobias, 2013. The wrong shape of insurance? What cross-sectional distributions tell us about models of consumption-smoothing. *Am. Econ. J. Macroecon.* 5, 107–140. <https://doi.org/10.1257/mac.5.4.107>.
- Doornik, Katherine, 2006. Relational contracting in partnerships. *J. Econ. Manag. Strategy* 15 (2), 517–548. <https://doi.org/10.1111/j.1530-9134.2006.00109.x>.
- Fong, Yuk-fai, Li, Jin, 2017. Relational contracts, limited liability, and employment dynamics. *J. Econ. Theory* 169 (C), 270–293. <https://doi.org/10.1016/j.jet.2017.02.006>.
- Foss, Sergey, Shneer, Vsevolod, Thomas, Jonathan P., Worrall, Tim, 2018. Stochastic stability of monotone economies in regenerative environments. *J. Econ. Theory* 173, 334–360. <https://doi.org/10.1016/j.jet.2017.11.004>.
- Hall, Robert E., 2005. Employment fluctuations with equilibrium wage stickiness. *Am. Econ. Rev.* 95 (1), 50–65. <https://doi.org/10.1257/0002828053828482>.
- Halonen, Maija, 2002. Reputation and the allocation of ownership. *Econ. J.* 112 (481), 539–558. <https://doi.org/10.1111/1468-0297.00729>.
- Kocherlakota, Narayana R., 1996. Implications of efficient risk sharing without commitment. *Rev. Econ. Stud.* 63 (4), 595–610. <https://doi.org/10.2307/2297795>.
- Kovrijnykh, Natalia, 2013. Debt contracts with partial commitment. *Am. Econ. Rev.* 103 (7), 2848–2874. <https://doi.org/10.1257/aer.103.7.2848>.
- Levin, Jonathan, 2003. Relational incentive contracts. *Am. Econ. Rev.* 93 (3), 837–857. <https://doi.org/10.1257/000282803322157115>.
- Ligon, Ethan, Thomas, Jonathan P., Worrall, Tim, 2002. Informal insurance arrangements with limited commitment: theory and evidence from village economies. *Rev. Econ. Stud.* 69 (1), 209–244. <https://doi.org/10.1111/1467-937X.00204>.
- Macleod, W. Bentley, Malcomson, James, 1989. Implicit contracts, incentive compatibility and involuntary unemployment. *Econometrica* 57 (2), 447–480. <https://doi.org/10.2307/1912562>.
- Ray, Debraj, 2002. The time structure of self-enforcing agreements. *Econometrica* 70 (2), 547–582. <https://doi.org/10.1111/1468-0262.00295>.
- Rayo, Luis, 2007. Relational incentives and moral hazard in teams. *Rev. Econ. Stud.* 74 (3), 937–963. <https://doi.org/10.1111/j.1467-937X.2007.00443.x>.
- Sigouin, Christian, 2003. Investment decisions, financial flows, and self-enforcing contracts. *Int. Econ. Rev.* 44 (4), 1359–1382. <https://doi.org/10.1111/1468-2354.t01-1-00113>.
- Thomas, Jonathan P., Worrall, Tim, 1988. Self-enforcing wage contracts. *Rev. Econ. Stud.* 55 (4), 541–554. <https://doi.org/10.2307/2297404>.
- Thomas, Jonathan P., Worrall, Tim, 1994. Foreign direct investment and the risk of expropriation. *Rev. Econ. Stud.* 61 (1), 81–108. <https://doi.org/10.2307/2297878>.