Working Paper Series

A Crises-Bailouts Game

WP 22-01

Bruno Salcedo
Western University

Bruno Sultanum
Federal Reserve Bank of Richmond

Ruilin Zhou
The Pennsylvania State University

FEDERAL RESERVE BANK OF RICHMOND®
Richmond ▪ Baltimore ▪ Charlotte

A Crises-Bailouts Game*

Bruno Salcedo[†] Bruno Sultanum[‡] Ruilin Zhou[§]

January 5, 2022

This paper studies the optimal design of a liability-sharing arrangement as an infinitely repeated game. We construct a noncooperative model with two agents: one active and one passive. The active agent can take a costly and unobservable action to reduce the incidence of crisis, but a crisis is costly for both agents. When a crisis occurs, each agent decides unilaterally how much to contribute mitigating it. For the one-shot game, when the avoidance cost is too high relative to the expected loss of crisis for the active agent, the first-best is not achievable. That is, the active agent cannot be induced to put in effort to minimize the incidence of crisis in a static game. We show that with the same stage-game environment, the first-best cannot be implemented as a perfect public equilibrium (PPE) of the infinitely repeated game either. Instead, at any constrained efficient PPE, the active agent "shirks" infinitely often, and when crisis happens, the active agent is "bailed out" infinitely often. The frequencies of crisis and bailout are endogenously determined in equilibrium. The welfare optimal equilibrium being characterized by recurrent crises and bailouts is consistent with historical episodes of financial crises, which features varying frequency and varied external responses for troubled institutions and countries in the real world. We explore some comparative statics of the PPEs of the repeated game numerically.

JEL classification: C73 · D82.

Keywords: Bailouts · Moral hazard · Repeated games · Imperfect monitoring · Second best.

*We thank Ed Green, V. Bhaskar, Neil Wallace, and Rishabh Kirpalani for helpful discussion and comments. The views expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Board of Governors.

[†]Department of Economics, Western University, bsalcedo@uwo.ca

[‡]Federal Reserve Bank of Richmond, bruno@sultanum.com.

[§]Department of Economics, The Pennsylvania State University, rzhou@psu.edu

1. Introduction

Some institutional arrangements intrinsically have high incentive costs. These costs surface in adverse situations. Sometimes attempts to lessen the damage may appear less cogent, while other times much harsher crisis management solutions are implemented. In the past, these different treatments seemed ad hoc and random, which economists often refer to as time inconsistent. In this paper, we argue that this seemingly random pattern of crisis management may be approximately optimal to sustain the relationship and to minimize long-term cost.

Historically, in episodes of financial turmoil, some troubled institutions have been bailed out, and others have not. As a result, the fate of these troubled parties ranges from complete failure/bankruptcy to full recovery.[1] Typically, a troubled institution gets bailed out on the ground that the alternative (failure) would have been a lot more costly, at least in the short run, since it might impose a big negative externality on many related parties. When this happens, economists are always quick to point out a fatal flaw of such rescue operation: the moral hazard—the "too big to fail" justification of bailout encourages behaviors that may lead to more failure. This incentive cost serves as the rationale for not bailing out some troubled institutions. We want to build an alternative theory on crisis/bailout that accounts for this very wide spectrum of outcomes. In view of the diverse outcomes, and of the tension between ex-ante and ex-post efficiency, questions about economic efficiency are especially salient.

We model this problem as a dynamic game between the "crisis-inflicting" party and the potential "help-to-clean-up" party. We construct a noncooperative, two-player model where an active agent takes costly unobservable action to reduce the incidence of crisis (avoidance). Whenever a crisis occurs, both parties suffer. Each agent decides unilaterally how much to contribute to contain the loss (mitigation). It is assumed both players have nontransferable utility. They can contribute directly to reduce the

[1]For example, while the majority of US companies sink or float on their own, the US government has consistently bailed out large corporations in the automobile industry such as GM and Chrysler. Among financial institutions, the Federal Reserve Bank significantly assisted large banks like Citigroup and Bank of America with loans and guarantees, while it let other large financial institutions such as Lehman Brothers and Washington Mutual fail during the 2007–2009 Great Recession. Among sovereign countries, the US government helped Mexico survive the 1994 Tequila crisis, while many countries suffered huge losses during the 1997 Asian financial crisis with little help from the IMF. During the recent Euro-zone crisis, Greece, Italy, Spain, and other potential problem countries received multiple rounds with varying magnitudes of bailouts.

loss from the crisis but not to increase each other's consumption. This assumption rules out direct subsidy from the passive agent to the active agent to pay his avoidance cost. The one-shot game can have any combination of avoidance/mitigation patterns as the static Nash equilibrium. In particular, when the avoidance cost is too high relative to the expected loss of crisis for the active agent, he cannot be induced to take the socially desirable but costly avoidance action at a static Nash equilibrium. As a result, crisis is more likely to happen. In this context, we consider what can be accomplished with the infinite repetition of the one-shot game.

We study the perfect public equilibria (PPE) of the repeated game. We show that in the environment where the active agent shirking is the only static Nash equilibrium of the stage game, the first-best outcome that requires that he takes the avoidance action every period cannot be implemented as a PPE of the infinitely repeated game. This is because in order to induce the active agent to take the costly avoidance action, the expected mitigation cost for him in case of crisis must be even higher. With both high avoidance cost and high mitigation cost, the active agent is better off doing nothing. In order to compensate the active player for taking the avoidance action *sometimes*, he has to be allowed to shirk other times and/or be bailed out (pay less than his share of the mitigation cost) when crises happen. Based on this intuition, we show that at any constrained efficient PPE, the active agent shirks infinitely often, and when crises happen, the active agent is bailed out infinitely often. As a result, crisis occurs more often compared to the first-best outcome.

Given that the constrained efficient allocation is necessarily achieved with stochastic shirking and bailout, we approximate the constrained efficient allocation with equilibrium allocation of finite-state automaton representation of the original game. With numerical examples, we show that a particular PPE, where the active agent shirks sometimes and is bailed out other times, can yield a welfare level much higher than the repetition of static Nash equilibrium, and the welfare loss relative to the allocation of the first-best varies with the parameters of the model. The corresponding PPE characterized by recurrent crises and bailouts is consistent with historical episodes of financial crises with varying frequency and varied external responses for troubled institutions and countries in the real world. As in Green and Porter (1984), such a phenomenon reflects an equilibrium that passes recurrently through several distinct states, rather than independent randomization by individual agents. It is a deliberate arrangement of using both occasional shirking and bailout as mechanisms

to incentivize good behavior of the active agent as often as possible. We explore some comparative statics of the PPEs of the N-state automaton numerically.

Our paper is closely related to the studies of the incentive problem induced by bailout policies. Most papers in the literature take particular institutional design and market structures very seriously but abstract from strategic dynamic interactions. We simplify the environment in several dimensions to make it tractable to illustrate the mechanism of stochastic bailout, but complicate the analysis by taking seriously the dynamic game played by the two parties involved in a bailout (the one who bails out and the one who is bailed out). This difference is not only technical, but also has interesting economic implications. In most papers, bailouts generate bad incentives to private agents. This is also true in our model if we only consider the stage game. However, we show that, in the repeated setting, promises of future bailouts are used to generate good incentives and reduce the incidence of crisis. Such strategic behavior is likely to be present in repeated interactions between long-lived large agents, such as members of the European Union or a government and a large corporation.

Two examples of papers highlighting the negative incentives generated by bailout policies are Farhi and Tirole (2012) and Chari and Kehoe (2015).[2] Chari and Kehoe (2015) study the time inconsistency problem of bailouts. The paper focus on the dynamic policy decision of a *bailout authority* who cannot commit to future actions (like the two players in our model). Farhi and Tirole (2012) consider a commitment problem from the government side, but they focus on the strategic complementarity of risk taking behavior from firms. The later paper study a finite stage game and the former assumes bailout policies to be noncontingent in agents' identities. As a result, neither studies the dynamic strategic behavior between a government and a firm.

Green (2010) and Keister (2016) highlight that bailouts can be welfare enhancing but not through incentives. Keister (2016) studies a version of Diamond and Dybvig (1983) that allows the government to divert tax funds from public goods to bailout banks and highlight two important implications of bailouts. On the one hand, bailouts induce bad behavior for banks, leading them to become less cautious and more illiquid. On the other hand, bailouts in their environment also provide insurance to depositors. Keister (2016) shows that, when the probability of a crisis is small, the insurance effect dominates. Similarly, in Green (2010), the welfare enhancing benefit of bailouts comes from the fact that, once we are in a regime with limited-liability firms, bailouts

[2]Other examples are Schneider and Tornell (2004) and Ennis and Keister (2009).

are necessary for firms to provide perfect risk sharing. The mechanism that makes bailouts welfare enhancing in these models, however, is very different from ours and doesn't have the incentive properties we highlight.

The paper is structured as follows. Section 2 introduces the stage game, while Section 3 describes the repeated game. Section 4 contains our main theoretical results. Section 5 uses numerical exercises to illustrate how the incentive mechanism works and to explore some comparative statics. Section 6 provides a discussion of alternative mechanisms under different assumptions. Finally, Section 7 concludes.

2. The stage game

There are two agents, agent 1 and agent 2, and two subperiods. In the first subperiod, agent 1 either takes an avoidance action to avert a crisis, $a = 1$, or not, $a = 0$. The cost of taking the avoidance action is $d > 0$, and the cost of not taking the avoidance action is normalized to zero. Agent 1's action a is unobservable to agent 2.

In the second subperiod one of two things happens: either there is a crisis, denoted by $\xi = 1$, or there is no crisis, denoted by $\xi = 0$. The probability of a crisis, conditional on agent 1's action in the first subperiod, $a \in \{0, 1\}$, is $\pi^a \in (0, 1)$. We assume that $\pi^1 < \pi^0$ so that taking the avoidance action reduces the probability of crisis. But agent 2 cannot infer agent 1's action from observing whether there is a crisis. Throughout the text, we refer to agent 1 as the active agent given that his action affects the probability of a crisis, and agent 2 as the passive one since he is forced to face the consequence of a crisis but has no influence on its occurrence.

In the event of a crisis, $\xi = 1$, the two agents can jointly mitigate the crisis. Let $m_i \geq 0$ denote agent i's contribution to mitigation. The crisis is mitigated if the total contribution of the two agents, $m_1 + m_2$, is no less than one. If the crisis is mitigated ($m_1 + m_2 \geq 1$), the cost to agent i is only his mitigation contribution m_i. If the crisis is not mitigated ($m_1 + m_2 < 1$), agent i suffers a loss $c_i > 0$ due to the crisis and his contribution m_i. If there is no crisis, $\xi = 0$, nothing needs to be mitigated and agents do not suffer any loss. It is implicit in the payoff structure that there is no transferable utility; contribution $m_1 + m_2$ is made only to mitigate a crisis. Neither party can consume it once it is made. The two agents can not make payment to each other in

the first subperiod either. In Section 6, we show that relaxing this assumption would greatly reduce the difficulty of achieving a better allocation in equilibrium. Figure 1 summarizes the structure of the stage game.

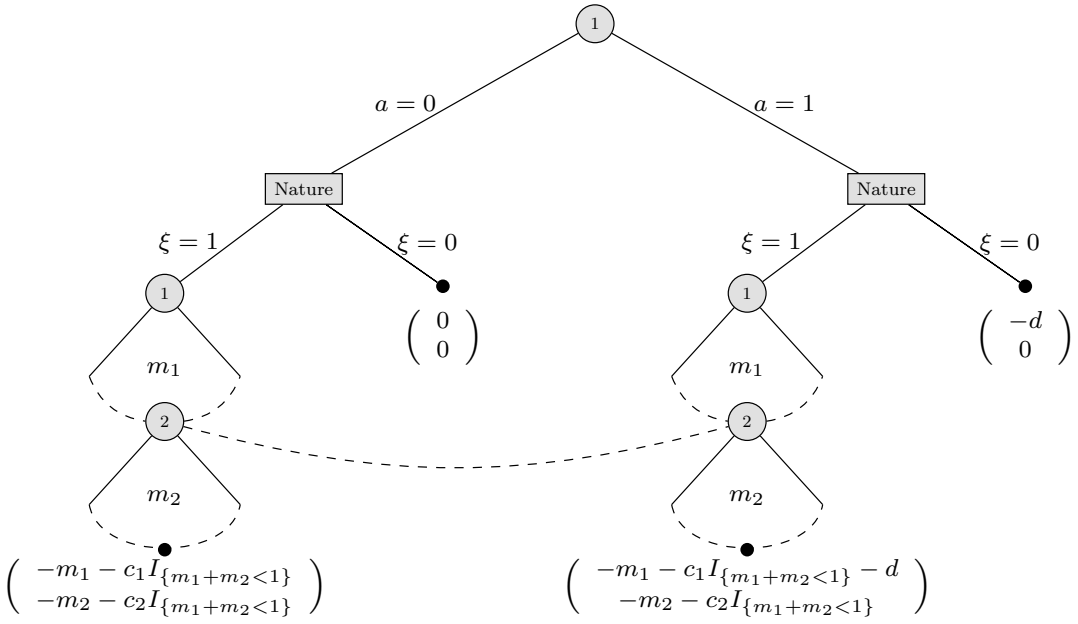

Figure 1: The stage game

We are interested in studying the case where mitigation after a crisis is ex-post efficient. Therefore, we make the following assumptions on the model parameters.

Assumption 1. For $i = 1, 2,\ c_i \in (0, 1)$, and $c_1 + c_2 > 1$.

Assumption 1 implies that neither agent alone is willing to mitigate the crisis, but together they should. Since the total cost of mitigation is less than the total loss if the crisis is not mitigated—that is, $1 < c_1 + c_2$—mitigation is efficient.

2.1. Equilibrium of the stage game

The structure of the game allows us to restrict attention to pure and public strategies without loss of generality. We thus solve for (pure-strategy, public-perfect) Nash equilibria of the two-subperiod normal-form game.

Denote the strategy for the active agent 1 by (a, m_1), and for the passive agent 2 by m_2, where a is agent 1's avoidance action, and m_i is agent $i = 1, 2$ mitigation

contribution. Given the strategy profile (a, m_1, m_2), agent 1 expected payoff is

$$u_1(a, m_1, m_2) = -ad - \pi^a(m_1 + c_1 I_{\{m_1+m_2\geq 1\}}), \tag{1}$$

and agent 2 expected payoff is

$$u_2(a, m_1, m_2) = -\pi^a(m_2 + c_2 I_{\{m_1+m_2\geq 1\}}). \tag{2}$$

As usual, a Nash equilibrium of the stage game is a strategy profile (a, m_1, m_2) such that (a, m_1) is a best response for agent 1 given the strategy m_2 for agent 2, and m_2 is a best response for agent 2 given the (a, m_1) for agent 1. The stage game has many equilibria, and it depends on parameters. We are interested in the parameter regions where the efficient outcome cannot be supported as equilibrium outcome.

First, there are multiple best responses in the mitigation stage after a crisis. In particular, by Assumption 1, after a crisis, not contributing to mitigation, $m_i = 0$, is always a best response if the other agent is doing the same—although this is ex-post inefficient. However, if agent 1 contributes m_1, and agent 2 contributes the remaining $1-m_1$, it must be $m_1 \leq c_1$ and $m_2 = 1-m_1 \leq c_2$. That is, for $(m_1, 1-m_1)$ to be both agent's best mitigation response to each other, we must have that $1 - c_2 \leq m_1 \leq c_1$. This mitigation outcome is ex-post efficient. The case with $m_1+m_2 > 1$ or $m_1+m_2 < 1$ with either $m_1 > 0$ or $m_2 > 0$ can be easily ruled out since at least one agent would be strictly better off by decreasing their mitigation contribution.

When deciding whether to take the avoidance action, agent 1 weighs the cost d against the expected gain of taking the action and successfully avoiding a crisis, which is either his expected contribution $(\pi^0 - \pi^1)m_1$ or expected loss $(\pi^0 - \pi^1)c_1$. Let $\hat{d} \equiv d/(\pi^0 - \pi^1)$ be the cost of avoidance adjusted by its impact on the probability. If $\hat{d} \leq c_1$, the efficient equilibrium where agent 1 takes the avoidance action and crisis is mitigated is always an equilibrium, and dominates all other equilibria. This is the uninteresting case since the static Nash equilibrium—or the repetition of it in the repeated game—achieves the first-best outcome. We ruled out this case and assume that $\hat{d} > c_1$.

Given the equilibrium restriction $m_1 \in [1 - c_2, c_1]$, the assumption $\hat{d} > c_1$ implies that $a = 0$ is always agent 1's optimal action. In this case, agent 1 never takes the avoidance action, and there are only two types of equilibrium: a nonmitigation

equilibrium where $(a, m_1, m_2) = (0, 0, 0)$, and a continuum of mitigation equilibria where $(a, m_1, m_2) = (0, m_1, 1 - m_1)$ and $m_1 \in [1 - c_2, c_1]$. These equilibria are inefficient when $\hat{d} < 1$ since the cost of action d is less than the expected social gain of avoiding a crisis $\pi^0 - \pi^1$. This is exactly the situation we are looking for. The inequality $\hat{d} > c_1$ guarantees that there is no avoidance in equilibrium, $\hat{d} < 1$ guarantees that and all the equilibria of the stage game is inefficient.

Assumption 2. $c_1 < \hat{d} < 1$.

We impose Assumption 2 on the parameters of the model, and in the next section, we investigate what can be achieved in this region for the repeated game.

3. The repeated game

In the *repeated game*, time is discrete and is indexed by $t \in \{1, 2, \ldots\}$. The two agents live forever and discount future payoffs with the same discount factor $\delta \in (0, 1)$. At the beginning of each period t, agents observe a payoff-irrelevant public signal $\theta_t \sim \mathcal{U}[0, 1]$, which is i.i.d. across periods. After observing the public signal, the agents play the stage game described in the previous section. The public signal allows agents to take correlated actions in each period. This serves a technical purpose—it convexifies the payoff set without explicitly considering randomized strategy for agent 1's action a_t. The public information at the beginning of date t is denoted by $h_t \in H_t$. It consists of the realization of all past and current public signals, the history of all past crises, and the history of all past contributions.

We focus on perfect Bayesian equilibria where both agents play *pure* and *public* strategies. Proposition A.2 establishes that this restriction is without loss of generality (see Appendix A for details). A public strategy for the active agent 1 is a sequence of measurable functions $\sigma_1 = (\alpha_t, \mu_{1t})_{t=1}^{\infty}$, where $\alpha_t(h_t) \in \{0, 1\}$ is the avoidance action of the active agent in period t given the public information h_t, and $\mu_{1t}(h_t) \geq 0$ is his contribution for the mitigation in case of a crisis. A public strategy for the passive agent is a sequence of functions $\sigma_2 = (\mu_{2t})_{t=1}^{\infty}$, where $\mu_{2t}(h_t) \geq 0$ is his contribution for the mitigation in case of a crisis. We denote a public strategy profile by $\sigma = (\sigma_1, \sigma_2)$.

The expected discounted utility for agent i from date t onward, given the strategy

profile σ and public history h_t is

$$v_{it}(\sigma, h_t) = (1-\delta)\mathbb{E}\left[\sum_{\tau=t}^{\infty} \delta^{\tau-t} u_i\Big(\alpha_\tau(h_\tau), \mu_{1\tau}(h_\tau), \mu_{2t}(h_\tau)\Big) \,\Big|\, h_t\right], \tag{3}$$

where the expectation is taken with respect to the crisis realizations from period t onward and the public signal from period $t+1$ onward. With slight abuse of notation, the average expected discounted utility for agent i at the beginning of the game is denoted by $v_i(\sigma) = \mathbb{E}[v_{i1}(\sigma, h_1)]$.

Definition 1. A public strategy profile σ^* is a *perfect public equilibrium* (PPE) if and only if $v_{it}(\sigma^*, h_t) \geq v_{it}(\sigma_i', \sigma_{-i}^*, h_t)$ for any agent i, any public strategy σ_i' any period $t \geq 1$, and any public history h_t.

We denote the set of PPE payoffs by $\mathcal{V}^* = \{v(\sigma^*) \mid \sigma^* \text{ is a PPE}\}$. A PPE always exists because unconditional repetition of a static Nash equilibrium of the stage game is a PPE. In Appendix A, we show that the set of PPE can be characterized recursively, using a modified version of the standard APS recursive decomposition, and we use this decomposition to establish some useful technical properties.

4. Optimal level of crises and bailouts

Under Assumptions 1 and 2, the first-best requires that in every period the active agent takes the avoidance action, and both agents mitigate a crises if it happens. However, the first-best is not achievable in a equilibrium of the stage game because the active agent never takes the avoidance action in any of them. In this section, we study how, and to what extent, welfare can be improved in the repeated setting. Our first finding is that, even in the repeated setting, the first-best cannot be attained. This result is a strong impossibility result because it holds for any discount factor of agents.

Given that the first-best is not achievable, we then turn our attention to constrained-efficient allocations by investigating the properties of PPEs that are Pareto efficient. In any Pareto efficient PPE, crises are always mitigated amd the welfare loss arises from the avoidance action not being taken every period. Furthermore, the frequency

of avoidance action depends on the agents' discount factor. For low discount factors, the active agent never takes the avoidance action in any equilibrium. Once the discount factor is greater than some threshold, the active agent takes the avoidance action infinitely often in any Pareto efficient PPE. Moreover, the passive agent has to *bailout* the active agent also infinitely often, where *"bailout"* has a precise sense we describe further ahead. The optimal frequency of avoidance and bailouts is determined endogenously.

4.1. The impossibility of implementing the first-best

Suppose that in some equilibrium the active agent takes the avoidance action with positive probability. The expected discounted payoff for the active agent at that moment (v_1) is a convex combination of his expected discounted payoff conditional on the event of a crisis (w_1^1), and his expected discounted payoff if there is no crisis (w_1^0). Also, because taking the avoidance action is costly, it must be the case that w_1^0 is strictly greater than w_1^1, so that the active agent finds it optimal to incur the cost. Moreover, w_1^1 cannot be too negative because of individual rationality. Using these facts, we show in the appendix that there is a fixed positive constant γ, such that $w_1^0 > v_1 + \gamma$. That is, whenever the active agent takes the avoidance action as part of a PPE and there is no crisis, his continuation value must increase by at least a fixed amount. Therefore, if there are no crises for a sufficiently long time interval, the implied continuation value required for the active agent to be willing to take the avoidance action value stops being feasible.[3] We thus obtain the following result.

Proposition 4.1. *There is no PPE in which the active agent takes the avoidance action almost surely at every period along the equilibrium path.*

4.2. Efficient mitigation

It is not possible to have avoidance played on every period. However, except for low discount factors, a PPE exists in which the active agent sometimes takes the avoidance action (see Lemma B.4 in the appendix). This requires the passive agent

[3]It is crucial for this proposition that the avoidance action is not observable, see Section 6.1. Hence, this is a result of moral hazard and not of the structure of the payoffs.

to provide incentives, for instance, by punishing the active agent after a crisis, or rewarding him if there is no crisis. Two possible ways to punish the active agent after a crisis are to let him suffer the cost of the crisis (no mitigation), or to ask him to contribute more than necessary to mitigate the crisis (money burning). Our second result is that neither of these forms of punishment schemes are optimal. In every constrained efficient PPE, agents contribute exactly as much as needed to mitigate a crisis when it happens.

Proposition 4.2. *In any constrained-efficient PPE, crises are efficiently mitigated, that is,* $\mu_{1t}(h_t) + \mu_{2t}(h_t) = 1$ *almost surely along the equilibrium path.*

This result is very natural since both of these forms of punishment are ex-post inefficient. However, the proof is far from trivial because, given that there is imperfect monitoring, some degree of inefficiency ex-post could be necessary to generate incentives ex-ante. This is a common feature of models with imperfect monitoring that can be traced back to Green and Porter (1984). We obtain the result because we show that there are always better ways to punish the active agent. As it turns out, any incentive scheme that can be generated in equilibrium via no-mitigation or money burning can also be generated by adjusting the shares of the mitigation cost in the future without incurring any efficiency losses due to either insufficient or excessive mitigation. The details of the proof are in Appendix B.

4.3. Bailouts as an incentive mechanism for avoidance

The difficulty in inducing the active agent to take the avoidance action is that the cost is too high for him to pay it on his own. The solution seems to be that the passive agent should help pay part of it. In a world with perfectly transferable utility, we could consider schemes where the passive agent directly subsidizes the active agent.[4] However, we have assumed that the agents' contributions can only be used to mitigate crises. In our environment, the only way for the passive agent to compensate the active agent is by sometimes paying more in mitigation cost after a crisis has occurred. When this happens, we call it a bailout.

[4]In Section 6.2, we study an extension with perfectly transferable utility, and show that the first-best can be achieved when both agents are sufficiently patient.

Definition 2. A *bailout* is a situation where a crisis occurs, the agents jointly contribute sufficient resources to mitigate it, and the contribution of the active agent is less than his private crisis loss, i.e., $\mu_{1t}(h_t) + \mu_{2t}(h_t) \geq 1$, and $\mu_{1t}(h_t) < c_1$.

We can show that bailouts are the only form of compensation available, and if the active agent is not compensated, then he has no reason to choose avoidance. It follows that bailouts are not only sufficient to induce the avoidance action, but also necessary.

Proposition 4.3. *In any PPE where the avoidance action is taken with positive probability, bailouts occur with positive probability.*

Proposition 4.3 shows that bailouts are necessary in order to support avoidance actions, but it says nothing about sufficiency nor about efficiency. When is it possible to support any avoidance at all? When is it efficient to do so? If the active agent expects to be bailed out in the future as a form of compensation, he may be willing to take the avoidance action, at least in some instances, and such arrangement is necessary for efficiency when feasible. The following proposition formalizes these results.

Proposition 4.4. *There exists $\tilde{\delta} \in (0, 1)$ such that:*

1. *If $\delta < \tilde{\delta}$, then every PPE (and therefore every constrained-efficient PPE) has avoidance played with probability zero at all periods.*
2. *If $\delta > \tilde{\delta}$, then in every constrained-efficient PPE the avoidance action is played infinitely often, and bailouts take place infinitely often.*

Proposition 4.4 indicates that, for low discount factors, it is *not* possible to induce the active agent to take any avoidance actions, and the set of efficient PPE essentially reduces to repetition of static Nash equilibria of the stage game. For higher discount factors, avoidance is not only possible, but it is also necessary for constrained efficiency. Propositions 4.1 and 4.4 combined imply that, when $\delta > \tilde{\delta}$, in *any* constrained-efficient PPE the active agent takes the avoidance action infinitely often, takes the nonavoidance action infinitely often, and is bailed out infinitely often.

To prove this fact, the key step is to show that having at least some avoidance is a Pareto improvement whenever it is incentive compatible. Formally, Lemma B.2 in the appendix asserts that, if it is possible to play avoidance at least once in *some*

PPE, then every PPE without avoidance is Pareto dominated by a PPE with avoidance. Because constrained Pareto efficiency requires continuation strategies to also be constrained-efficient, this implies than, whenever possible, it is optimal to have avoidance infinitely often. Because of Proposition 4.3, doing so requires also having bailouts infinitely often.

5. Automata: endogenously determined frequencies of crises and bailouts

From previous sections, we learned that, since the active agent's private incentive does not align with the social one, i.e., $c_1 < \hat{d} < 1$, the way to align incentives is by bailing out the active agent infinitely often. But how bailouts work as a mechanism to generate incentives? How good is it? Is it close to the first-best? And how does it change with the primitives of the model, such as, the cost to avert crisis, d, the private costs of agents, c_1 and c_2, and the effectiveness of the avoidance action, $\pi_0 - \pi_1$? In this section, we use numerical methods to investigate these questions.

We approximate the second-best by considering PPEs where equilibrium behavior can be described by finite-state automata. An automaton consists of four components: states, an initial distribution over states, a transition rule, and a mapping from states to actions. Let Ω be a finite set of states. States are mapped into actions by $\alpha_1 : \Omega \to A_i$ and $\mu_i : \Omega \times X \to M$, $i = 1, 2$. Given the current state ω_t, the active agent takes the action $\alpha_1(\omega_t)$, and, after a crisis state ξ_t is realized, the agents' contributions are given by $\mu_i(\omega_t, \xi_t)$. After actions are realized, the state ω_{t+1} for the next period is randomly drawn according to the transition rule $\eta : \Omega \times X \times M \times M \to \Delta(\Omega)$. The initial state for period 1 is drawn according to the initial distribution $\eta_0 \in \Delta(\Omega)$.

An automaton describes a profile of public strategies for the repeated game. In fact, if we didn't restrict attention to *finite* automata, every profile of public strategies could be described by an automaton (Mailath and Samuelson, 2006, pp. 230). However, for computational reasons, in all of our numerical exercises, we restrict attention to finite automata with a fixed upper bound on the number of states in Ω. In what follows, we provide an example in the form of an automaton, where avoidance takes place at some but not at other times.

5.1. An illustrative example of equilibrium mechanism

Here, we illustrate how bailouts can be used to induce avoidance action and, thus, improve efficiency. Consider the set of parameters

$$\delta = 0.95, \quad \pi^1 = 0.2, \quad \pi^0 = 0.9, \quad d = 0.5, \quad c_1 = 0.6, \quad c_2 = 0.5.$$

With this set of parameters, the adjusted avoidance cost is higher than the cost for agent 1 and avoidance is socially efficient,

$$c_1 = 0.6 < \hat{d} \approx 0.7143 < 1.$$

The first-best has an expected total cost of 0.7, which is not obtainable in any PPE by Proposition 4.1. The total cost that can be obtained in a static Nash equilibrium is 0.9, which implies a welfare loss (relative to the first-best) of 28.57 percent.

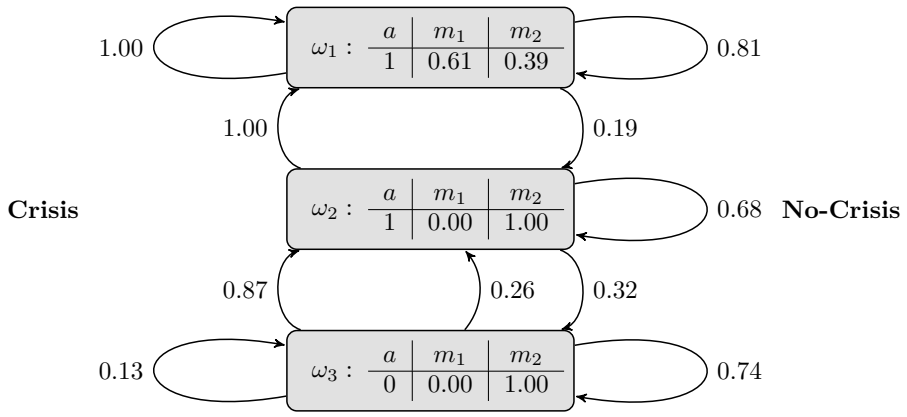

Figure 2: Automaton

Figure 2 describes the PPE that minimizes the total expected long-run cost among all PPEs that can be described by a four-state automata, with one state being the minmax equilibrium. The equilibrium works as follows. Agents start in state ω_1, where the strategy profile is $(a, m_1, m_2) = (1, 0.61, 0.39)$. In this state, agent 1 is supposed to take the avoidance action, but his private cost in the crisis alone does not generate incentives to do so since $m_1 = 0.61 < \hat{d} \approx 0.7143$. As a result, in order to generate incentives, when there is no crisis, the state switches to ω_2 with probability

0.19. The ω_2 is a bailout state since $m_1 = 0.00 < c_1$. The "reward" of a bailout in the future helps generating incentives for the avoidance action. That is, the probability of going to this bailout state, compensates for the fact that $m_1 = 0.61 < \hat{d}$, aligning the private and social incentives to take the avoidance action. In ω_2, the strategy profile is $(1, 0.00, 1.00)$. Again, agent 1 is supposed to take the avoidance action, but now he has even less incentives to do so since his contribution to mitigation is now zero. This time, to generate incentives, the equilibrium moves to state ω_3 with probability 0.32 if there is no crisis. The state ω_3 has an even stronger form of bailout because the active agent contribution in mitigating crisis is zero, and he takes no avoidance action.[5] There is a fourth state ω_4, which is not in the figure, with strategy profile of nonavoidance/no-mitigation (the minmax equilibrium). This state is out of the equilibrium path and works as a punishment state in case of a detectable deviation.

Table 1: Summary statistics of the PPE

State ω	Invariant distribution	$u_1(\omega)$	$u_2(\omega)$	$V_1(\omega)$	$V_2(\omega)$	Welfare	Welfare loss (%)
ω_1	0.50	-0.622	-0.078	-0.539	-0.177	-0.748	2.23
ω_2	0.38	-0.500	-0.200	-0.478	-0.250	-0.727	3.87
ω_3	0.12	-0.000	-0.900	-0.420	-0.328	-0.716	6.86
LRA	–	-0.501	-0.222	-0.501	-0.222	-0.724	3.41

Notes: LRA refers to the long-run averages, which correspond to the expected values evaluated using the invariant distribution. $u_i(\omega)$ denotes agent i's expected payoff for the period when the state is ω. $V_i(\omega)$ denotes agent i's total discounted expected payoff when the state is ω.

Although this automaton PPE is not in the Pareto frontier, it provides a lower bound on what can be achieved by a constrained efficient allocation. Table 1 provides summary statistics of the equilibrium. In state ω_1, the expected discounted total cost is 0.748, which is only 2.23 percent greater than the minimum feasible one 0.7. On average, crisis occurs 28.4 percent of the time, compared to 20 percent at the first-best. When crisis does happen, bailout occurs 70.3 percent of the time. But the striking result is that, even though the passive agent bailout the active agent over 70 percent of the crisis, the expected present value of his cost is only 0.17. For a comparison, in the best equilibrium for the passive agent with no bailouts, the expected present value of his cost is 0.36. That is, by optimally choosing a bailout policy, the passive agent can reduce his cost with crisis by half.

[5]The probability of moving to a state preferred by agent 1 is always higher when there is no crisis. Hence, the automata is reminiscent of the revision strategies used in Rubinstein and Yaari (1983) and Radner (1985).

5.2. Comparative statics

The equilibrium displayed in Figure 2 illustrates how bailouts can be used in order to induce avoidance in equilibrium. In this subsection, we study how properties of this equilibrium change with key parameters of the model: the avoidance cost, d, the private costs of nonmitigated crisis, (c_1, c_2), and the probabilities of crises, (π^0, π^1). For each set of parameters, we found the PPE that minimizes the total discounted long-run cost among six states automata. Then, we compare the implied long-run probabilities of avoidance, crisis and bailouts, as well as the long-run average cost of avoidance and the agents' mitigation payments.

The impact of changes in the avoidance cost — Table 2 displays features of the equilibrium outcome for different values of the avoidance cost d. As one could expect, when the avoidance cost increases, avoidance action is taken less frequently and, therefore, crisis happens more often. The average avoidance cost (column 5) is nonmonotone, reflecting the more costly avoidance action is taken less often. The incidence of bailouts (column 4) is nonmonotone, similar to agent 1's mitigation cost (column 6). These changes reflect the structure of the equilibrium. Bailouts are the mechanism where agent 2 compensates agent 1 for bearing the avoidance cost alone. As d increases, the compensation needed to generate incentives for the avoidance action also increases.

Table 2: The impact of changes in the avoidance cost[1]

d	$\mathbb{P}(a=1)$	$\mathbb{P}(\xi=1)$	$\mathbb{P}(m_1 < c_1)$	$\mathbb{E}(d)$	$\mathbb{E}(m_1)$	$\mathbb{E}(m_2)$	Expected Total Cost (%)[2]
0.45	0.9569	0.2302	0.4647	0.4306	0.0784	0.1518	101.66
0.50	0.8571	0.3000	0.7937	0.4285	0.0340	0.2660	104.08
0.55	0.7598	0.3682	0.8942	0.4179	0.0182	0.3500	104.81
0.60	0.7115	0.4019	0.9301	0.4269	0.0143	0.3877	103.61
0.65	0.6851	0.4204	0.8727	0.4453	0.0227	0.3977	101.85

Note: The probabilities and expectations are evaluated using the implied invariant distribution.
[1] Other parameters are $\delta = 0.9$, $\pi^1 = 0.2$, $\pi^0 = 0.9$, $c_1 = 0.6$ and $c_2 = 0.5$.
[2] Expressed as percentage of the first-best expected total cost.

The impact of changes in the private costs of a crisis — Table 3 displays features of the equilibrium outcome for different values of (c_1, c_2). When c_1 and c_2 increases, both agents' minmax payoff decreases. The impact of c_1 on the equilibrium outcome is substantial. Increasing c_1 from 0.55 to 0.65 reduces the long-run average cost from about 106.4 percent to 101.8 percent of the first-best; the incidence of crisis is reduced

by approximately one-third (from 0.36 to 0.24); and bailouts are reduced to 60 percent from 89 percent. The reason is that increasing c_1 helps aligning agent 1 private cost of a crisis, c_1, with the social cost of a crisis, which is the mitigation cost 1. Agent 2 private cost of crisis, c_2, has little effect on the equilibrium outcome since he is a passive agent and has no private information.

Table 3: The impact of changes in the private costs of a crisis[1]

c_1	c_2	$\mathbb{P}(a=1)$	$\mathbb{P}(\xi=1)$	$\mathbb{P}(m_1<c_1)$	$\mathbb{E}(d)$	$\mathbb{E}(m_1)$	$\mathbb{E}(m_2)$	Expected Total Cost (%)[2]
0.55	0.5	0.7763	0.3566	0.8870	0.3882	0.0177	0.3388	106.39
	0.7	0.7729	0.3590	0.8875	0.3864	0.0176	0.3414	106.49
0.60	0.5	0.8571	0.3000	0.7937	0.4285	0.0340	0.2660	104.08
	0.7	0.8529	0.3029	0.7119	0.4265	0.0391	0.2638	104.20
0.65	0.5	0.9363	0.2446	0.6136	0.4681	0.0655	0.1791	101.82
	0.7	0.9349	0.2456	0.5420	0.4675	0.0722	0.1733	101.86

Note: The probabilities and expectations are evaluated using the implied invariant distribution.
[1] Other parameters are $\delta = 0.9$, $\pi^1 = 0.2$, $\pi^0 = 0.9$, and $d = 0.5$.
[2] Expressed as percentage of the first-best expected total cost.

The impact of changes in crisis probabilities — Table 4 displays the long-run expected total cost above the first-best for different combinations of (π^0, π^1). The other parameters are set to $\delta = 0.95$, $d = 0.5$, $c_1 = 0.6$, and $c_2 = 0.5$. The cells with symbol "$-$" represent the cases where the parameters do not satisfy Assumption 2. The effect of (π^0, π^1) on the welfare cost is not uniform. Combinations of (π^0, π^1), with $\hat{d}$ either closer to c_1 or 1, lead to lower cost. This means that sometimes decreasing π^0 reduces the welfare cost, while sometimes increasing π^0 reduces the welfare cost. The same is true for π^1.

On the other hand, for combinations of (π^0, π^1) with the same $\hat{d}$ (that is, $\pi^0 - \pi^1$ constant), higher π^0 and π^1 always lead to a lower welfare cost. The interpretation of these results is not simple. One could think that higher π^1 means that avoidance is less effective in preventing crisis, which could imply a higher cost, but this is not true. The correct measure is $\pi^0 - \pi^1$, how much the probability of crisis decreases by the avoidance action. With $\pi^0 - \pi^1$ held constant, the only impact is increasing π^0. Higher π^0 implies that the minmax utility of agent 1 is lower, hence, it is easier to generate incentives for avoidance.

Table 4: Total cost above the first-best (%)[1]

π^1 \ π^0	0.65	0.70	0.75	0.80	0.85	0.90	0.95
0.10	3.25	5.71	8.49	9.05	8.00	2.94	–
0.15	–	2.70	4.65	6.46	6.59	4.41	1.28
0.20	–	–	2.20	3.68	4.67	4.54	2.74
0.25	–	–	–	1.75	2.89	3.30	2.89
0.30	–	–	–	–	1.38	2.19	2.26

[1] Long-run average as percentage of the first-best.

The impact of changes in agents' discount factor — Table 5 displays features of the equilibrium outcome for different values of the discount rate δ. As one could expect, lower δ is associated with lower welfare. Increasing δ from 0.6 to 0.9 decreases the total cost in about 4 percent of the first-best. This pattern reflects that when δ is high, agents are more willing to cooperate since they care more about punishments in the future.

Table 5: The impact of changes in agents' discount factor[1]

δ	$\mathbb{P}(a=1)$	$\mathbb{P}(\xi=1)$	$\mathbb{P}(m_1<c_1)$	$\mathbb{E}(d)$	$\mathbb{E}(m_1)$	$\mathbb{E}(m_2)$	Expected Total Cost (%)[2]
0.6	0.7111	0.4022	0.5978	0.3555	0.1063	0.2960	108.25
0.7	0.7630	0.3659	0.8345	0.3815	0.0636	0.3023	106.77
0.8	0.7979	0.3415	0.7491	0.3989	0.0444	0.2970	105.78
0.9	0.8571	0.3000	0.7937	0.4285	0.0340	0.2660	104.08

Note: The probabilities and expectations are evaluated using the implied invariant distribution.
[1] Other parameters are $\pi^1 = 0.2$, $\pi^0 = 0.9$, $c_1 = 0.6$, $c_2 = 0.5$, and $d = 0.5$.
[2] Expressed as percentage of the first-best expected total cost.

6. Alternative mechanisms

We have shown that the first-best cannot be achieved as a PPE of the repeated game, and that whenever avoidance is possible in equilibrium, every constrained efficient PPE involves bailouts infinitely often. In this section, we consider two alternative mechanisms that can help to improve welfare. We analyze one model where the avoidance action is perfectly observed, and one where the passive agent can directly subsidize the active agent. In both cases, it is still the case that either bailouts or direct transfers are necessary for the active agent to take the avoidance action.

However, unlike our benchmark model, these alternative specifications admit PPE that achieve the first-best when agents are patient enough.

6.1. The avoidance action is observable

In our benchmark model, we assumed that the avoidance action of the active agent is private. The passive agent could only make imperfect inferences about it via the realization of crises. Now, consider the alternative specification where a is perfectly observable to both agents. This allows agents to use strategy profiles that bail out the active agent if and only if he takes the avoidance action, but it does not change the fact that bailouts are necessary for avoidance.

Proposition 6.1. *In any PPE of the game with observable actions, if the avoidance action is taken with positive probability, then bailouts occur with positive probability.*

To illustrate the difference from the unobservable action case, consider the following simple strategy profile. Along the equilibrium path, the active agent always takes the avoidance action, and crisis is always mitigated.

$$(\alpha_t(h_t), \mu_{1t}(h_t), \mu_{2t}(h_t)) = (1, m_1^*, 1 - m_1^*),$$

for some fixed constant $m_1^* > 0$, which is specified ahead. After any deviation, then the active agent chooses $a = 0$ forever after, and both agents never again make positive mitigation contributions. We show in the appendix that, if the discount factor is sufficiently high, then one such grim trigger strategy exists, which is a PPE. Since there is always avoidance and mitigation along the equilibrium path, this strategy profile implements the first-best.

Proposition 6.2. *There exists $\tilde{\delta}' \in (0, 1)$ such that, if $\delta > \tilde{\delta}'$, then the game with observable actions admits a PPE where the active agent takes the avoidance action at every period and after every history.*

6.2. Monetary transfers

The previous analysis depends crucially on the assumption of nontransferable utility. That is, if agent 1 takes the avoidance action, he has to pay the cost d by himself.

Moreover, both agents' contributions for cleanup can only be used to mitigate crises. Suppose we relax this assumption by allowing the passive agent to directly transfer resources to agent 2 for consumption. More precisely, suppose that at any date t and after any history h_t, agent 2 can make a transfer $\beta_{t1}(t, h_t) \geq 0$ if there is a crisis and a transfer $\beta_{t0}(t, h_t) \geq 0$ if there is no crisis. These transfers enter the stage game payoffs as an additive term. That is, the stage-game payoffs for the active (passive) agent in the game with transfers are exactly those from the game without transfers plus (minus) whatever transfer he receives (makes).

A version of Proposition 4.3 continues to hold in this modified model. For the active agent to be willing to take the avoidance action, he must expect some form of compensation. The only difference is that the passive agent has new forms of compensation available. Agent 2 can still compensate agent 1 by bailing him out, by contributing sufficient resources so that the cost incurred by agent 1 in case of a crisis is less than c_1. Additionally, agent 2 can transfer resources to agent 1 in periods where there are no crises. Any equilibrium with avoidance must involve at least one of these forms of compensation.

Proposition 6.3. *In any PPE of the game with transfers where the active agent takes the avoidance action with positive probability, agent* 2 *compensates agent* 1 *by having either* $\beta_{t0}(h_t) > 0$ *or* $\beta_{t1}(h_t) - \mu_{1t}(h_t) > -c_1$*, or both with positive probability.*

To illustrate the difference from the nontransferable-utility case, consider the following simple strategy profile for the game with transfers.

$$(\alpha_t(h_t), \mu_{1t}(h_t), \mu_{2t}(h_t)) = (1, m_1^*, 1 - m_1^*),$$

and

$$(\beta_{t0}(h_t), \beta_{1t}(h_t)) = (b^*, 0),$$

for all t and every h_t along the equilibrium path, where $m_1^* \in (0, 1)$ and $b^* > 0$ are fixed constants specified in Appendix B.7. That is, agent 1 always takes the avoidance action and contributes m_1^* when there is a crisis, and agent 2 compensates agent 1 with b^* units of consumption when there is no crisis. The transfer b^* can be viewed as a subsidy to agent 1 from agent 2 in no-crisis time. In case of a detectable deviation, the agents switch to play the one-shot Nash equilibrium with no-avoidance and no

mitigation forever. We show in the appendix that, if the discount factor is high enough, this strategy profile constitutes a PPE of the game with transfers. Hence, when the agents are patient enough, the first-best is attainable in equilibrium.

Proposition 6.4. *There exists $\tilde{\delta}'' \in (0,1)$ such that, if $\delta > \tilde{\delta}''$, then the game with transfers admits a PPE where the active agent takes the avoidance action at every period and after every history.*

This subsidy scheme is simple theoretically but may not be easy to implement in reality. For example, it might be difficult to justify paying Greece's government every period—subsidy in normal time and mitigation in crisis time—to the public!

7. Conclusion

We have studied a liability-sharing problem between two asymmetric parties in an infinitely repeated game. The main frictions in the model are unobserved action by the active player (moral hazard) and nontransferable utility between the two parties. With this model, we want to make several points. First, there are environments where, conditional on some social arrangement (such as the European Monetary Union) having already been formed to share some risk, shirking and bailouts are not only consistent with equilibrium behavior, but also necessary to achieve constrained optimal. The incentive cost may be too high to ask for outcomes devoid of these *vices*. Insisting otherwise is unrealistic. The high incentive cost to the social arrangement should be considered before any coalition/arrangement is made rather than ex-post trying to eliminate it. Second, stochastic shirking and bailout may be necessary features of the approximately efficient outcome. Roughly speaking, when the active player is expected to shirk, there is no need for incentive to induce his current-period effort, and hence bailout is likely as a reward from the passive player to the active player for future effort. In a period when the active player is supposed to put in effort, he is unlikely to be bailed out. The constrained optimal requires coordination between the two parties. This coordination can be accomplished with the use of n-state automata and the correlated equilibrium given the automata. To achieve the constrained optimal, the fine-tuning tools for incentive provision are the levels of mitigation contribution and the transition probability from any state to any other

state conditional on current state and outcome. These transition probabilities are endogenously chosen, unlike the exogenous sunspot type of modeling device. Third, our numerical simulation results show that the equilibrium of the n-state automata, with n optimally chosen (approximate second-best), can achieve quite high level of welfare relative to the first-best. Fourth, if one thinks that the nontransferable utility assumption is too strong, relaxing the assumption can improve welfare of the two parties, but it will not eliminate "bailout." With transferable utility, the payment from the passive player 2 to the active player 1 is simply shifted from ex-post (after the crisis happens) to ex-ante (before the effort is exerted), but does not disappear.

The model is very schematic: it does not have any realistic features such as different maturities of debt instrument, sovereign default, renegotiation of debt, yield, fiscal and monetary policies, etc. This is intentional and meant to illustrate the mechanism of stochastic bailout as an incentive device.

Appendix

A. Recoursive analysis of the set of PPE

With agents playing public strategies only, the repeated game has a recursive structure. After an arbitrary history, the continuation strategy profile of a PPE is an equilibrium profile of the original game. The standard way to characterize the set of PPE values is to use the self-generation procedure introduced in Abreu, Pearce, and Stacchetti (1990) (APS). This appendix establishes an analogous procedure and shows that our restriction to pure and public strategies is without loss of generality.[6]

A.1. Incentive constraints

We begin by providing three necessary conditions that any PPE must satisfy after each history: one feasibility condition and two incentive constraints. Consider any PPE of the repeated game $\sigma^* = (\alpha^*, \mu_1^*, \mu_2^*)$, and any arbitrary public history $h_t \in H_t$. Let $s^* = (a^*, m_1^*, m_2^*)$ denote the action profile dictated by σ^* for period t given h_t, i.e., $(a^*, m_1^*, m_2^*) = (\alpha_t^*(h_t), m_{1t}^*(h_t), m_{2t}^*(h_t))$. Also, let $w^* = (w_i^*(\xi))_{i=1,2;\xi=0,1} \in \mathbb{R}^{2\times 2}$ denote the profile of continuation expected average discounted values from date $(t+1)$ onward given σ^* as a function of the crisis state on date t, i.e., $w_i^*(\xi) := \mathbb{E}\Big[v_{it+1}(\sigma^*, h_{t+1})|h_t, \xi_t = \xi\Big]$. With this notation, note that we can write the following *feasibility condition*

$$v_{it}(\sigma^*, h_t) = g_i(s^*, w^*), \tag{F}$$

where $g_i : (\{0,1\} \times \mathbb{R}_+^2) \times \mathbb{R}^{2\times 2} \to \mathbb{R}$ is the function given by

$$g\Big((a, m_1, m_2), w\Big) = (1-\delta)u_i\Big(a, m_1, m_2\Big) + \delta\Big[\pi^a w_i(1) + (1-\pi^a)w_i(0)\Big]. \tag{4}$$

There are two kinds of necessary date-t conditions for σ^* to be a PPT. The first

[6] We cannot simply apply the procedure from APS because our model differs from theirs in the monitoring structure, our stage game is a multistage game, and we allow for public randomization but exclude individual mixed strategies.

condition refers to the mitigation contributions. If a crisis were to arise on t, each agent i could unilaterally decide to contribute exactly the minimum amount required to mitigate it, Agent i's cost from the crisis would be $-\max\{0, 1-m^*_{-i}\}$. Alternatively, if $m^*_{-i} < 1$, agent i could decide to *not* contribute anything to mitigate the crisis, and incur a cost of $-c_i$. By doing so, agent i's ex-post cost due to the crisis on the period would be $-\min\{c_i, \max\{0, 1-m^*_{-i}\}\}$, and his continuation value would be no worse than his minimax $-\pi^0 c_i$. For σ^* to be a PPE, this potential deviation cannot be strictly profitable. That is, it must be the case that

$$\begin{aligned}(1-\delta)k_i(m_1^*, m_2^*) + \delta w_i^*(1) &\geq -(1-\delta)\min\{c_i, \max\{0, 1-m^*_{-i}\}\} - \delta\pi^0 c_i \\ &\geq -(1-\delta+\delta\pi^0)c_i \end{aligned} \tag{M}$$

where $k_i : \mathbb{R}^2_+ \to \mathbb{R}_+$ is agent i's cost from a date-t crisis as a function of the mitigation contributions, i.e.,

$$k_i(m_1, m_2) = -m_i - c_i I_{\{m_1+m_2<1\}}. \tag{5}$$

We call Condition (M) the *mitigation constraint* for agent i.

Secondly, suppose that $\alpha^* = 1$. The active agent—agent 1—could deviate at period t by *not* taking the avoidance action and following σ_1^* after that. Since this deviation is not observable, agent 1 would expect in equilibrium that agent 2 would continue to follow σ_2^*. For this deviation to not be profitable, the expected discounted utility for the active agent in case there is no crisis ($\delta w_1^*(0)$ should be greater than if there is a crisis ($(1-\delta)k_1(m_1^*, m_2^*) + \delta w_1^*(1)$). Moreover, it should be sufficiently greater to compensate the active agent for the private cost of taking the avoidance action, that is,

$$(1-\delta)\hat{d} \leq \delta w_1^*(0) - \Big[(1-\delta)k_1(m_1^*, m_2^*) + \delta w_1^*(1)\Big].$$

If $\alpha^* = 0$, the converse inequality must hold. After doing some simple algebra, we can summarize both cases via the following *avoidance constraint*

$$0 \leq (-1)^{a^*}\left[\hat{d} + k_1(m_1^*, m_2^*) + \frac{\delta}{1-\delta}\Big(w_1^*(1) - w_1^*(0)\Big)\right]. \tag{A}$$

A.2. APS decomposition

So far, we have argued that conditions (F), (M), and (A) are necessary for a strategy profile to be a PPE. In what follows, we will show that they are also sufficient to characterize the set of PPE payoffs $\mathcal{V}^*$. In particular, we will show that a vector of feasible payoffs v is attainable in equilibrium if and only if it can be attained by (a distribution over)[7] action profiles and continuation values that satisfy such conditions. To formalize this idea, we make use of the following definition.

Definition 3. Given an action profile $s^* = (a^*, m_1^*, m_2^*)$, a profile of continuation values $w^* = (w_i(\xi))_{i=1,2;\xi=0,1}$, and a set $\mathcal{V} \subset \mathbb{R}^2$, the pair (s^*, w^*) is said to be *admissible* with respect to $\mathcal{V}$ if and only if: (a) it satisfies the mitigation constraints (M) for $i = 1, 2$, (b) it satisfies the admissibility constraint (A), and (c) $w^*(\xi) = (w_1^*(\xi), w_2^*(\xi)) \in \mathcal{V}$ for $\xi = 0, 1$.

In our setting with public randomization, the relevant self-generating operator is the one introduced by Cronshaw and Luenberger (1994). For any set $\mathcal{V} \subseteq \mathbb{R}^2$ and every action profile $s = (a, m_1, m_2)$, define:

$$B_s(\mathcal{V}) = \left\{ v \in \mathcal{V} \,\middle|\, \exists w \text{ such that } (s, w) \text{ is admissible w.r.t. } W \text{ and } v = g(s, w) \right\}. \quad (6)$$

Intuitively, $B_s(\mathcal{V})$ would be the set of payoff profiles that could be obtained by playing s on the first period and using continuation values from the set $\mathcal{V}$, in such a way that there are no profitable one-shot deviations on the first period. To take into account the possibility of public randomization, let

$$B(\mathcal{V}) = \text{co}\left(\bigcup_{s \in S} B_s(\mathcal{V}) \right). \quad (7)$$

A set $\mathcal{V} \subseteq \mathbb{R}^2$ is said to be *self-generating* if and only if $\mathcal{V} \subseteq B(\mathcal{V})$. The following proposition states that, the set of PPE satisfies some desirable properties and, using this notion of self-generation, the following APS-like result applies to our setting.

[7] So far in this section, we have not yet discussed public randomization. The action profile (a^*, m_1^*, m_2^*) specifies the pure actions chosen after θ_t is realized, and the continuation values w^* are taken to be the average continuation values integrating over θ_{t+1}. Public randomization enters implicitly in the convex hull operation in (7).

Proposition A.1 (APS decomposition)**.** *The set of PPE payoffs $\mathcal{V}^*$ is the largest self-generating set, and $B^n(\mathcal{V}) \to \mathcal{V}^*$ for every bounded set $\mathcal{V} \subseteq \mathbb{R}^2_+$ such that $\mathcal{V}^* \subseteq \mathcal{V}$.*

Proposition A.1 serves two purposes. First, it enables us to use a recursive approach to characterize the set of PPE. We use this approach implicitly in the proofs of our main results. Secondly, it allows us to establish some desirable properties for the set of PPE, which are summarized in the following proposition.

Proposition A.2. *The set of PPE payoffs $\mathcal{V}^*$ is nonempty, compact and convex, and is increasing with respect to the discount factor δ. Moreover, $\mathcal{V}^*$ would remain unchanged if we allowed player 1 to use private strategies and we allowed both players to use mixed strategies.*

A.3. Proofs of recursive characterization

Proving propositions A.1 and A.2 requires a number of technical lemmas. Because many of the proof steps are standard, we omit some details and refer the interested reader to Mailath and Samuelson (2006) instead.

Lemma A.3 (One-shot deviation principle)**.** *An individually rational strategy profile is a PPE if and only if it admits no profitable one-shot deviations.*

Proof. This lemma is analogous to Proposition 2.2.1 in (Mailath and Samuelson, 2006, pp 25), and the corresponding arguments can be easily adapted to our setting. The structure of the argument is as follows. Fix an individually rational strategy profile. Suppose that there is a profitable deviation, and let v be the difference in values between the proposed strategy profile and the profitable deviation. Since the set of feasible individually rational payoffs of the stage-game is bounded, we know that there is some number T such that the payoffs after T periods amount to less than $\|v\|/2$. Thus, there must also be a profitable deviation of length at most T. If the deviation in the last period is profitable, then the proof is complete. If not, then there is a profitable deviation of length at most $T-1$ periods. By induction, this implies that there is a profitable deviation of length 1. ■

Lemma A.4 (Self-generation)**.** *If $\mathcal{V} \subseteq \mathbb{R}^2$ is bounded and self-generating, then $\mathcal{V} \subseteq \mathcal{V}^*$.*

Proof. Take any point $v \in \mathcal{V}$. Since $v \in \mathcal{V} \subseteq \mathcal{V}^*$, Carateheodory's theorem implies that the value profiles $b_v^1, b_v^2, b_v^3 \in \cup_s B_s(\mathcal{V})$ exist, and a vector of weights $(\lambda_v^1, \lambda_v^2, \lambda_v^3) \in \Delta^3$ such that $v = \sum_{n=1}^3 \lambda_v^n b_v^n$. For each b_v^n, since $b_v^n \in \cup_s B_s(\mathcal{V})$, there exist s_v^n and a profile of continuation values w_v^n such that $b_v^n = g(s_v^n, w_v^n)$, and (s_v^n, w_v^n) is admissible w.r.t. $\mathcal{V}$.

Now, fix some $v^* \in \mathcal{V}$. We will construct a PPE σ^* such that $v^* = v(\sigma^*)$. For that purpose, we will construct a sequence of (public) history-dependent continuation values $v_t^0 : H_t \to \mathcal{V}$ such that $v_t^0(h_t)$ does *not* depend on θ_t. Along the equilibrium path, σ^* is defined as a function of v^0:

$$\sigma_t^*(h_t) = \begin{cases} s^1_{v_t^0(h_t)} & \text{if } \theta_t \leq \lambda^1_{v_t^0(h_t)} \\ s^2_{v_t^0(h_t)} & \text{if } \lambda^1_{v_t^0(h_t)} < \theta_t \leq \lambda^2_{v_t^0(h_t)} \\ s^3_{v_t^0(h_t)} & \text{if } \theta_t > \lambda^2_{v_t^0(h_t)} \end{cases}.$$

Along the equilibrium path, v_t^0 is defined recursively with $v_1^0(h_1) \equiv v^*$ and:

$$v_{t+1}^0\left(h_{t+1}\right) = \begin{cases} w^1_{v_t^0(h_t)}(\xi_t) & \text{if } \theta_t \leq \lambda^1_{v_t^0(h_t)} \\ w^2_{v_t^0(h_t)}(\xi_t) & \text{if } \lambda^1_{v_t^0(h_t)} < \theta_t \leq \lambda^2_{v_t^0(h_t)} \\ w^3_{v_t^0(h_t)}(\xi_t) & \text{if } \theta_t > \lambda^2_{v_t^0(h_t)} \end{cases}.$$

After any observable deviation from σ^*, agents turn to the autartic equilibrium with $\sigma_t^*(h_t) = (0,0,0)$ and $v_t^* = -\pi^0 c$ for any subsequent public history h_t. It is straightforward to see that v_t^0, and thus σ_t^*, are measurable. For every public history h_t along the equilibrium path, by construction we have that

$$\begin{aligned} v_t^0(h_t) &= \sum_{n=1}^3 \lambda^n_{v_t^0(h_t)} g\left(s^n_{v_t^0(h_t)}, w^n_{v_t^0(h_t)}\right) = \mathbb{E}_t\Big[(1-\delta)u(\sigma_t^*(h_t)) + \delta v_{t+1}^0(h_{t+1})\Big] \\ &= (1-\delta)\mathbb{E}_t\left[u(\sigma_t^*(h_t)) + \delta u(\sigma_{t+1}^*(h_{t+1})) + \frac{\delta^2}{1-\delta} v_{t+2}^0(h_{t+2})\right] \\ &= (1-\delta)\mathbb{E}_t\left[\sum_{\tau=0}^{\infty} \delta^\tau u(\sigma_{t+\tau}^*(h_{t+\tau}))\right] + \lim_{\tau\to\infty} \delta^\tau \mathbb{E}_t\left[v_{t+\tau}^0(h_{t+\tau})\right] = v_t\left(\sigma^*, h_t\right). \end{aligned}$$

Hence, we have that $v^* = v_1^0(h_1) = v(\sigma^*)$. Finally, since actions and continuation values are admissible at every period, we know that there are no profitable single deviations. (Recall that the conditions (A) and (M) that define admissibility are

precisely the requirement that there should be no one-shot deviations at the avoidance and mitigation stages, respectively). By Lemma A.3, this implies that σ^* is a PPE. ■

Lemma A.5 (Factorization)**.** $\mathcal{V}^*$ *is self-generating*

Proof. Fix an arbitrary point $v^* \in \mathcal{V}^*$, and let σ^* be the PPE that generates it. For each possible realization of the date-1 public signal $(\theta_1) \in [0,1]$, let $\sigma'|_{\theta_1}$ denote the corresponding continuation strategy from period t_1 onward (assuming that there are no detectable deviations), and let $w_{\theta_1} = v(\sigma'|_{\theta_1})$. Since σ^* is measurable, it follows that w_{θ_1} is also measurable. Since σ^* is a PPE, we know that $w_{\theta_1} \in \mathcal{V}^*$ for all θ_1. Moreover, Lemma A.3 implies that there are no profitable single-shot deviations from σ^* on period 1. Therefore, $(\sigma_1^*(\theta_1), w_{\theta_1})$ is admissible w.r.t. $\mathcal{V}^*$ for each realization of θ_1. Hence, $g\left(\sigma_1^*(\theta_1), w_{\theta_1}\right) \in B_{\sigma_1^*(\theta_1)}(\mathcal{V}^*)$. This implies that:

$$v^* = \int_0^1 g\left(\sigma_1^*(\theta_1), w_{\theta_1}\right)\, d\theta_1 \in \text{co}\left(\bigcup_s B_s(\mathcal{V}^*)\right) = B(\mathcal{V}^*),$$

thus completing the proof. ■

Lemma A.6. *If* $\mathcal{V}$ *is compact, then* $B(\mathcal{V})$ *is compact.*

Proof. Fix some $a \in \{0,1\}$. We will start by showing that $B_a(\mathcal{V}) := \cup_m B_{a,m}(\mathcal{V})$ is compact. Consider any sequence (v^n) in $B_a(\mathcal{V})$ converging to some $v^* \in \mathbb{R}^2$. By construction, sequences (m^n) and (w^n) exist such that $v^n = g(a, m^n, w^n)$, and (a, m^n, w^n) is admissible w.r.t. $\mathcal{V}$. Since it is contained in a compact space, the sequence (m^n, w^n) has a subsequence converging to some limit (m^*, w^*). Since $\mathcal{V}$ and $\mathbb{R}^2_+$ are closed, we know that $m^* \in \mathbb{R}^2_+$ and $w^* \in \mathcal{V}$. Since g is continuous, we know that $v^* = g(a, m^*, w^*)$. Since the incentive constraints (A) and (M) are defined by continuous functions, we know that (a, m^*, w^*) is admissible w.r.t. $\mathcal{V}$. Hence, $v^* \in B_a(\mathcal{V})$. Since this was for arbitrary convergent sequences, this means that $B_a(\mathcal{V})$ is closed.

Now, since the payoffs of the stage game are all nonpositive and $\mathcal{V}$ is bounded, then $B_a(\mathcal{V})$ is bounded above. Since admissibility implies that the values have to be conditionally individually rational, it is also bounded below. Hence, $B_a(\mathcal{V})$ is compact. Since a finite union of compact sets is compact, we have that $\cup_s B_s(\mathcal{V}) = \cup_a B_a(\mathcal{V})$ is compact. The result then follows from the fact that the convex hulls of compact sets are compact. ■

Proof of proposition A.1. Since B is $\subseteq$-monotone by construction, Lemma A.4 implies that $\mathcal{V}^*$ contains the union of all self-generating sets. By Lemma A.5, this implies that $\mathcal{V}^*$ is the largest self-generating set. Since $B(\mathcal{V})$ is convex for any $\mathcal{V}$ by construction, Lemma A.5 also implies that $\mathcal{V}^*$ is convex. Now, fix any bounded set $\mathcal{V}$ such that $\mathcal{V}^* \subseteq \mathcal{V}$. Let $\bar{\mathcal{V}}$ be the closure of $\mathcal{V}$, and define the sequence $\{\mathcal{V}^n\}_{n=1}^{\infty}$ by $\mathcal{V}^0 = \bar{\mathcal{V}}$ and $\mathcal{V}^{n+1} = B(\mathcal{V}^n)$ for $n = 1, 2, \ldots$. By definition of B and Lemma A.6, we know that $\mathcal{V}^n$ is a $\subseteq$-decreasing sequence of compact sets and therefore has a (Hausdorff) limit $\mathcal{V}^\infty = \cap_n \mathcal{V}^n$, and this limit is compact. Since B is $\subseteq$-monotone and $\mathcal{V}^*$ is self-generating, we know that $\mathcal{V}^* = B^n(\mathcal{V}^*) \subseteq B^n(\mathcal{V}^0) = \mathcal{V}^n$ for all n, and thus $\mathcal{V}^* \subseteq \mathcal{V}^\infty$.

It remains to show that $\mathcal{V}^\infty$ is self-generating. For this purpose, we combine the proofs from lemmas A.4 and A.6. Consider any $v^* \in \mathcal{V}^\infty$. By construction we know that $v^* \in B(\mathcal{V}^n)$ for all n. Therefore, sequences $\left\{(b^{nk}, \lambda^{nk}, s^{nk}, w^{nk})_{k=1}^3\right\}_{n=1}^{\infty}$ exist such that $v^* = \sum_{k=1}^3 \lambda^{nk} b^{nk}$, $b^{nk} = g(s^{nk}, w^{nk})$ and (s^{nk}, w^{nk}) is admissible w.r.t. $\mathcal{V}^n$ for all n. Since it is contained in a compact space, the sequence $\{b^{nk}, \lambda^{nk}, s^{nk}, w^{nk}(0), w^{nk}(1, m^{nk})\}$ has a subsequence converging to some limit $(b^{*k}, \lambda^{*k}, m^{*k}, w^{*k})$. Since all the relevant sets are closed, we know that the limit belongs to the set where we want it to be. Since g is continuous, we know that $b^{*k} = g(s^{*k}, w^{*k})$. Since the incentive constraints are defined by continuous functions, we know that (s^{*k}, w^{*k}) is admissible w.r.t. $\mathcal{V}^\infty$. It is straightforward to see that $v^* = \sum_{k=1}^3 \lambda^{*k} b^{*k} \in B(\mathcal{V}^\infty)$. Therefore $\mathcal{V}^\infty$ is self-generating and, by Lemma A.4, $\mathcal{V}^\infty \subseteq \mathcal{V}^*$. ■

We are now in a position to prove our claim about the restriction to pure-public strategies being without loss of generality. One could extend the definition of equilibrium in the obvious way to allow for mixed strategies that depend on private information. The following proposition states that the set of equilibrium payoffs would not change. The reason for this is because the new set would be self-generating in the original sense, and thus it would be contained in $\mathcal{V}$.

Proof of proposition A.2. $\mathcal{V}^*$ is nonempty because unconditional static repetition of a Nash equilibrium of the stage game constitutes a PPE. Compactness and convexity follow directly from the first part of the proof of Proposition A.1. For δ-monotonicity, it is easy to see from the definition of B that, if $\mathcal{V}^* \subseteq \mathcal{V}$, then $B(\mathcal{V})$ is $\subseteq$-monotone with respect to δ. Hence, the set of PPE payoffs is also $\subseteq$-monotone with respect to δ.

Finally, it remains to argue that the restriction to pure strategies is without loss of generality. The complete proof is technical and burdensome. Hence, we only present a sketch of the proof, but a formal proof can be provided upon request. The definitions of equilibrium and $v(\sigma)$ can be easily extended to allow for mixed and private strategies in the obvious way. Let $\tilde{\mathcal{V}}$ be the corresponding set of equilibrium payoffs with the modified definitions. Fix some $v^* \in \tilde{\mathcal{V}}$, and let σ^* be the (possibly mixed or private) strategy profile that generates it and constitutes an equilibrium. Now delegate all the randomization to θ_1, define continuation values in the obvious way, and show that the resulting pairs $(\sigma|\theta_1, w|\theta_1)$ are admissible in accordance with Definition 3 w.r.t. $\tilde{V}$. Intuitively, this occurs because $\mathbb{R}^2_+$ is convex and thus there is no need to randomize mitigation contributions. Moreover, m_1 and m_2 are chosen after observing a, and thus there is no need to randomize the avoidance action. This implies that $\tilde{\mathcal{V}}$ is self-generating and is thus contained in $\mathcal{V}^*$. ■

B. Proofs of the main results

B.1. Preliminaries

Throughout this section, we use the notation $\langle a, m_1, m_2 \rangle$ to denote the stationary strategy profile for the repeated game that consists of repeating (a, m_1, m_2) in every period and after any public history. From the analysis in Section 2, we know that $(0, m_1, 1 - m_1)$ is a NE of the stage game as long as $m_1 \in [1 - c_2, c_1]$. Therefore, $\langle 0, m_1, 1 - m_1 \rangle$ is a PPE as long as $m_1 \in [1 - c_2, c_1]$. We use this fact repeatedly in the subsequent proofs.

Each agent i can guarantee a minmax payoff of $-\pi^0 c_i$ by never making any positive contributions and, if $i = 1$, then never taking the avoidance action. Hence, every equilibrium payoff $v \in \mathcal{V}^*$ must satisfy the individual rationality conditions $v_i \geq -\pi^0 c_i$, for $i = 1, 2$. The set of feasible and individually rational payoffs corresponds to the shaded area in Figure 3. Each of the diagonal lines in the figure corresponds to the feasible payoffs that can be attained with efficient mitigation with and without taking the avoidance action, respectively. The thick blue line corresponds to the set of equilibrium payoffs that can be achieved by unconditional repetitions of static Nash

equilibria of the stage game.

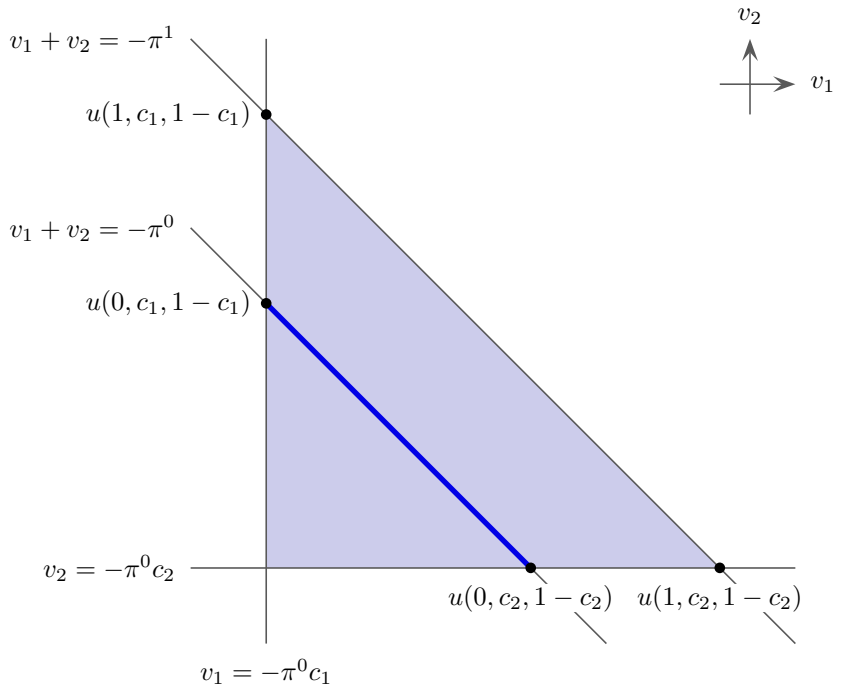

Figure 3: Feasible and individually rational payoffs, and stationary PPE payoffs.

B.2. Proof of Proposition 4.1

Let $\sigma^* = (\alpha^*, \mu_1^*, \mu_2^*)$ be a PPE and fix a history (h_t with $\alpha_t^*(h_t) = 1$, (if there are no such PPE and histories, then the proposition is true). For σ^* to be an equilibrium, it must satisfy the feasibility and incentive constraints from Section A.1. First, the feasibility condition (F) implies that

$$v_{1t}(\sigma^*, h_t) = -(1-\delta)d + \pi^1\Big((1-\delta)k_1(m_1^*, m_2^*) + \delta w_1^*(1)\Big) + (1-\pi^1)\delta w_1^*(0), \qquad (8)$$

where $m_i^* = \mu_{it}^*(h_t)$ is agent i's equilibrium mitigation contribution for date-t according to σ^*, and $w_i^*(\xi)$ are his equilibrium continuation values. The avoidance constraint

(A) can be written as

$$\Big((1-\delta)k_1(m_1^*, m_2^*) + \delta w_1^*(1)\Big) \leq \delta w_1^*(0) - (1-\delta)\hat{d}. \tag{9}$$

Combining this constraint with the mitigation constraint (M), we have that:

$$\begin{aligned} & \delta w_1^*(0) - (1-\delta)\hat{d} \geq -(1-\delta+\delta\pi^0)c_1. \\ \Rightarrow \quad & \delta\pi^0\hat{d} \geq -\delta w_1^*(0) - (1-\delta+\pi^0\delta)\Big(c_1 - \hat{d}\Big). \end{aligned} \tag{10}$$

Solving for $((1-\delta)k_1(m_1^*, m_2^*) + \delta w_1^*(1))$ in (8), substituting in (9), and rearranging terms yields:

$$\delta\pi^0\hat{d} \leq \left(\frac{\delta^2}{1-\delta}\right) w_1^*(0) - \left(\frac{\delta}{1-\delta}\right) v_{1t}(\sigma^*, h_t). \tag{11}$$

Combining (10) and (11) and doing some more algebra we obtain:

$$w_i^*(0) \geq v_{1t}(\sigma^*, h_t) + \gamma, \qquad \gamma := (1-\delta+\pi^0\delta)\left(\frac{1-\delta}{\delta}\right)\Big(\hat{d} - c_1\Big). \tag{12}$$

Hence, we have established that whenever the active agent chooses $a = 1$ in equilibrium and there is no crisis, his expected value must increase at least by a fixed factor γ. The assumption $\hat{d} > c_1$ guarantees that $\gamma > 0$. This implies that, if there is no crisis for n subsequent periods and the active agent keeps choosing $a = 1$ with probability 1, then we must have

$$w_i^n > v_{1t}(\sigma^*, h_t) + n\gamma,$$

where w_i^n is agent 1's expected discounted value at period $t + n$. Since the set of feasible payoffs of agent 1 is bounded above by 0, it must be the case that after a long enough history of no crisis, agent 1 takes the nonavoidance action. Otherwise, w_i^n would be greater than 0. On the other hand, since $\pi^1, \pi^0 \in (0, 1)$, any finite sequence of no crisis occurs with positive probability. Therefore, taking the avoidance action at every period with probability 1 cannot be part of a PPE. ■

B.3. Proof of Proposition 4.2

The central step of the proof of Proposition 4.2 is to establish Lemma B.1 below. The lemma can be understood as follows. In equilibrium, starting from a history where either crises are not mitigated ($m_1 + m_2 < 1$) or money is burned ($m_1 + m_2 > 1$), it is possible to make a Pareto improvement that makes agent 2 *strictly* better off, while keeping the payoff of agent 1 constant.

Lemma B.1. *If (s^0, w^0) is admissible w.r.t. the set of PPE $\mathcal{V}^*$ and $m_1 + m_2 \neq 1$, then $v' \in \mathcal{V}^*$ exists such that $v_1 = g_1(s^0, w^0)$ and $v_2 > g_2(s^0, w^0)$, where g_1 and g_2 are defined as in (4).*

Proof. There are three different cases to consider.

Case 1.— Suppose that crises are not mitigated, i.e., $m_1^0 + m_2^0 < 1$. Consider the alternative action profile $s' = (a', m_1', m_2')$ with $a' = a^0$, $m_1' = m_1^0 + c_1$, and $m_2' = \min\{0, 1 - m_1'\}$. Note that, $m_1' + m_2' = 1$, that is, crises are mitigated according to the new action profile. Hence, the cost in case of a crisis for the active agent remains unchanged, that is,

$$k_1(m_1', m_2') = -m_1' = -m_1^0 - c_1 = k_1(m_1^0, m_2^0),$$

where k_1 is the crisis cost function as defined in (5). In contrast, the cost in case of a crisis for the active agent goes down since

$$k_2(m_1', m_2') = -\min\{0, 1 - m_1'\} \geq -(1 - c_1) > -c_2 \geq k_2(m_1^0, m_2^0),$$

where the strict inequality follows from Assumption 1. This implies that $g_1(s', w^0) = g_1(s^0, w^0)$ and $g_2(s', w^0) > g_2(s^0, w^0)$. Moreover, since the mitigation contributions only enter the incentive constraints (A) and (M) through k_i, the pair (s', w^0) is admissible w.r.t. $\mathcal{V}^*$. Thus, by Proposition A.1, $g(s', w^0) \in \mathcal{V}^*$.

Case 2.— Suppose that the passive agent burns money in case of a crisis, i.e., $m_1^0 + m_2^0 > 1$ and $m_2 > 0$. Consider the alternative action profile $s' = (a', m_1', m_2')$ with $a' = a^0$, $m_1' = m_1^0$, and $m_2' = \min\{0, 1 - m_1'\}$. As in the previous case, we have $k_1(m_1', m_2') = k_1(m_1^0, m_2^0)$, and $k_2(m_1', m_2') < k_2(m_1^0, m_2^0)$. Moreover, this implies that $g(s', w^0) \in \mathcal{V}^*$, $g_1(s', w^0) = g_1(s^0, w^0)$ and $g_2(s', w^0) > g_2(s^0, w^0)$.

Case 3.— Suppose that *only* the active agent burns money in case of a crisis,

i.e., $m_1^0 > 1$ and $m_2^0 = 0$. We begin by showing that, in this case, the equilibrium continuation value for the active agent in case of a crisis cannot be the maximum equilibrium value, i.e., we must have $w_2^0(1) < \max\{v_2 \mid v \in \mathcal{V}^*\}$.

For that purpose, consider the alternative continuation value profile w' and the alternative action profile $s' = (a', m_1', m_2')$ with $a' = 1$, $m_1' = m_1^0$, $m_2' = 0$, and $w_i'(\xi) = w_i^0(1)$ for $i = 1, 2$ and $\xi = 0, 1$. Since $w_1'(1) = w_1'(0)$, the avoidance constraint (A) for (s', w') can be written as $\hat{d} \leq -k_1(m_1^0, m_2^0)$. It is satisfied because $-k_1(m_1^0, m_2^0) = m_1^0 > 1$, and Assumption 2 requires that $\hat{d} < 1$. The mitigation contributions and continuation values after a crisis are the same under (s', w') and (s^0, w^0). Hence, we know that (s', w') satisfies the mitigation constraints (M) for both agents. Hence, (s', w') is admissible w.r.t. $\mathcal{V}^*$ and, by Proposition A.1, $g(s', w') \in \mathcal{V}^*$.

Since $w_2'(0) = w_2'(1) = w_2^0(1)$, it follows that

$$g_2(s', w') = (1 - \delta)u_2(s') + \delta w_2^0(1) = \delta w_2^0(1).$$

Now, it is easy to show that there are no equilibria where agent 1 always mitigates crises on its own. This implies that $w_2^0(1) < 0$ and, therefore, $g_2(s', w') < w_2^0(1)$. Hence, $v^* \in \mathcal{V}^*$ exists such that $v_2^* > w_2^0(1)$.

Now, we can return to showing that (s^0, w^0) is inefficient. For each $\varepsilon \in (0, 1)$, consider the alternative continuation value profile w^ε and the alternative action profile $s^\varepsilon = (a^\varepsilon, m_1^\varepsilon, m_2^\varepsilon)$ with $a^\varepsilon = a^0$, $m_2^\varepsilon = 0$,

$$m_1^\varepsilon = m_1^0 + \frac{\delta}{1 - \delta}\varepsilon(v_1^* - w_1^0(1)),$$

$w_i^\varepsilon(0) = w_i^0(0)$, and

$$w_i^\varepsilon(1) = (1 - \varepsilon)w_i^0(1) + \varepsilon v_i^*,$$

for $i = 1, 2$.

Now, fix any ε sufficiently small so that $m_1^\varepsilon > 1$. Note that m_1^ε was chosen specifically so that

$$\begin{aligned} -(1 - \delta)m_1^\varepsilon + \delta w_1^\varepsilon(1) &= -(1 - \delta)\left[m_1^0 + \frac{\delta}{1 - \delta}\varepsilon(v_1^* - w_1^0(1))\right] + \delta\left[(1 - \varepsilon)w_i^0(1) + \varepsilon v_i^*\right] \\ &= -(1 - \delta)m_1^0 + \delta w_i^0(1). \end{aligned}$$

Hence, both the avoidance (A) and the mitigation constraint (M) for the active agent are satisfied by $(s^\varepsilon, w^\varepsilon)$ and $g_1(s^\varepsilon, w^\varepsilon) = g_1(s^0, w^0)$. As for the passive agent, since $k_2(m_1^\varepsilon, m_2^\varepsilon) = k_2(m_1^\varepsilon, m_2^\varepsilon)$, $w_2^\varepsilon(0) = w_2^0(0)$ and $w_2^\varepsilon(1) > w_2^0(1)$, we know that his mitigation constraint is satisfied by $(s^\varepsilon, w^\varepsilon)$, and $g_2(s^\varepsilon, w^\varepsilon) > g_2(s^0, w^0)$. Also, by Proposition A.2, we know that $\mathcal{V}^*$ is convex and thus $w^\varepsilon(1) = (w_1^\varepsilon(1), w_2^\varepsilon(1)) \in \mathcal{V}^*$. Hence, $(s^\varepsilon, w^\varepsilon)$ is admissible w.r.t. $\mathcal{V}^*$ and thus, by Proposition A.1, $g(s^\varepsilon, w^\varepsilon) \in \mathcal{V}^*$. ■

With Lemma B.1, it is easy to prove Proposition 4.2.

Proof of Proposition 4.2. Let σ^* be a PPE, and let $H^0 \subseteq \cup_{t=1}^\infty H_t$ be the (possibly empty) set of public histories h_t such that (a) $\mu_{1t}^*(h_t) + \mu_{2t}^*(h_t) \neq 1$ and (b) $\mu_{1t}^*(h_t') + \mu_{2t'}^*(h_t) = 1$ for every public history h_t', which is a predecessor of h_t. By Lemma B.1, we know that for every such history $h_t \in H^0$, a strategy profile σ^{h_t} exists such that $v_1(\sigma^{h_t}) = v_{1t}(\sigma^*, h_t)$ and $v_2(\sigma^{h_t}) > v_{2t}(\sigma^*, h_t)$. Let σ' be the strategy profile that mimics σ^* until it reaches a public history $h_t \in H^0$ and follows σ^{h_t} from then onward (treating h_t as the empty history). Since continuation values for the passive agent remain unchanged, and continuation values for the active agent only go up, it follows that σ' is a PPE. Finally, if H^0 is nonempty and is reached with positive probability, then $v_1(\sigma') = v_1(s^*)$ and $v_2(\sigma') > v_2(s^*)$. ■

B.4. Proof of Proposition 4.3

Proof. We will show that in a PPE where there are no bailouts, the active agent always takes the no-avoidance action almost surely. Consider a strategy profile with no bailouts, i.e., such that for almost every history (h_t), either $\mu_{1t}(h_t) \geq c_1$ or $\mu_{1t}(h_t) + \mu_{2t}(h_t) < 1$. This implies that $u_1(\sigma_t(h_t)) \leq -d - \pi^1 c_1$ for histories with $\alpha_t(h_t) = 1$, and $u_1(\sigma_t(h_t)) \leq -\pi^0 c_1$ for histories with $\alpha_t(h_t) = 0$. Assumption 2 implies that $-d - \pi^1 c_1 < -\pi^0 c_1$, and thus $v_1(\sigma) \leq -\pi^0 c_1$, with strict inequality whenever $\alpha_t(h_t) = 1$ for some history set of histories $\{h_t\}$, which is reached with positive probability along the equilibrium path. Individual rationality requires $v_1(\sigma) \geq -\pi^0 c_1$. Hence, a strategy profile with no bailouts can satisfy individual rationality only if $\alpha_t(h_t) = 0$ almost surely along the equilibrium path. ■

B.5. Proof of Proposition 4.4

The proof of Proposition 4.4 makes use of three lemmas. Lemma B.2 simply asserts that, whenever it is possible to have avoidance, it is possible to do it in a way that dominates the best static Nash equilibria of the stage game in terms of total cost.

Lemma B.2. *If it is possible for agent* 1 *to choose the avoidance action at least once in at least one PPE, then a PPE exists with total cost less than* π^0*, i.e., if a PPE* σ^* *and a public history* h_t *exist such that* $\alpha_t^*(h_t) = 1$*, then* $(v_1, v_2) \in \mathcal{V}^*$ *exists such that* $v_1 + v_2 > -\pi^0$*.*

Proof. Suppose it is possible for agent 1 to choose the avoidance action at least once in at least one PPE. Then, by Proposition A.1, there exist a profile of continuation values w^* and a profile of actions $s^* = (a^*, m_1^*, m_2^*)$ with $a^* = 1$ such that (s^*, w^*) is admissible w.r.t. $\mathcal{V}^*$. If either $w_1^*(0) + w_2^*(0) > -\pi^0$ or $w_1^*(1) + w_2^*(1) > -\pi^0$, then the proof is complete. Hence, for the rest of the proof, we assume that $w_1^*(0) + w_2^*(0) \leq -\pi^0$ and $w_1^*(1) + w_2^*(1) \leq -\pi^0$.

Individual rationality implies that $w_i^*(0) \geq -\pi^0 c_i$ for $i = 1, 2$. Hence, we know that $-\pi^0 c_1 \leq w_1^*(0) \leq -\pi^0(1 - c_2)$. This implies that $m_1^0 \in [1 - c_2, c_1]$ exists such that $w_1^*(0) = -\pi^0 m_1^0$ (See Figure 3). Now, consider the alternative continuation value profile w' and the alternative action profile $s' = (1, c_1, 1 - c_1)$, $w_1'(1) = -\pi^0 c_1$, $w_2'(1) = -\pi^0(1 - c_1)$, $w_1'(0) = -\pi^0 m_1^0$ and $w_2'(0) = -\pi^0(1 - m_1^0)$.

Since $\langle 0, c_1, 1 - c_1 \rangle$ and $\langle 0, m_1^0, 1 - m_1^0 \rangle$ is a PPE, we know that $w'(\xi) \in \mathcal{V}^*$ for $\xi = 0, 1$. Also, it is straightforward to verify that (s', w') satisfies the mitigation constraints (M) for both agents. In order to show that (s', w') satisfies the avoidance constraint (A), first consider the pair (s^*, w^*). Since (s^*, w^*) is admissible w.r.t. $\mathcal{V}^*$, it must satisfy the mitigation constraint

$$(1 - \delta) k_1(m_1^*, m_2^*) + \delta w_1^*(1) \geq -(1 - \delta + \delta\pi^0) c_1,$$

and the avoidance constraint

$$(1 - \delta) k_1(m_1, m_2^*) + \delta w_1^*(1) \leq \delta w_1^*(0) - (1 - \delta)\hat{d}.$$

Together, these two constrains imply that

$$-(1-\delta+\delta\pi^0)c_1 \leq \delta w_1^*(0) - (1-\delta)\hat{d},$$

which is precisely the avoidance constraint for (s', w').

We have shown that (s', w') is admissible w.r.t. $\mathcal{V}^*$. By Proposition A.1, this implies that $g(s', w') \in \mathcal{V}^*$. Finally, note that

$$g_1(s', w') + g_2(s', w') = (1-\delta)(-\pi^1) - \delta\pi^0 < -\pi^0,$$

where the last inequality follows from Assumption 2. ■

Lemma B.3, is the crucial step of the proof. It states that if it is possible to have avoidance then, starting from any nonavoidance PPE, it is possible to improve the payoff of the passive agent without affecting the payoff of the active agent. Since continuation values of efficient PPE must always lay on the upper boundary of $\mathcal{V}^*$, this implies that avoidance must happen infinitely often. By Proposition 4.3, this implies that bailouts must happen infinitely often as well.

Lemma B.3. *If a PPE exists with total cost less than π^0, then every PPE without avoidance is Pareto dominated by a different PPE that increases the payoff of agent 2 while keeping the payoff of agent 1 unchanged, i.e., for every PPE σ^0 such that $\alpha_t^0(h_t) = 0$ almost surely, $v' \in \mathcal{V}^*$ exists such that $v_1' = v_1(\sigma^0)$ and $v_2' > v_2(\sigma^0)$.*

Proof. Further below we will show that $v^1, v^2 \in \mathcal{V}^*$ exist such that $v_i^i = -\pi^0 c_i$, and $v_{-i}^i > -\pi^0(1-c_i)$, for $i = 1, 2$. For now, take this fact as given, and let σ^0 be PPE without avoidance. Since $\alpha_t^0(h_t) = 0$ for all h_t, we know that $v_1(\sigma^0) + v_2(\sigma^0) \leq -\pi^0$. Moreover, individual rationality implies that $-\pi^0 c_1 \leq v_1(\sigma^0) \leq -\pi^0(1-c_2)$. Therefore, some $\mu \in (0,1)$ exists such that $v_1(\sigma^0) = \mu v_1^1 + (1-\mu)v_1^2$. Let $v^\mu = \mu v^1 + (1-\mu)v^2$. By construction, we know that $v_1^\mu + v_2^\mu > -\pi_0$, which implies that $v_2^\mu > v_2(\sigma^0)$, see Figure 4. The result then follows because $\mathcal{V}^*$ is convex (Proposition A.2) and, consequently, $v^\mu \in \mathcal{V}^*$.

It only remains to show the existence of v^1 and v^2. We will only show existence of v^1. The proof for v^2 is analogous. Let $v^* \in \mathcal{V}^*$ be a PPE with total expected discounted cost less than π^0, i.e. such that $v_1^* + v_2^* > -\pi^0$. Individual rationality implies that $v_1^* \geq -\pi^0 c_1$ If $v_1^* = -\pi^0 c_1$, then we can simply set $v^1 = v^*$. For the rest of the proof, we consider the case that $v_1^* = -\pi^0 c_1 + \Delta$ for some $\Delta > 0$.

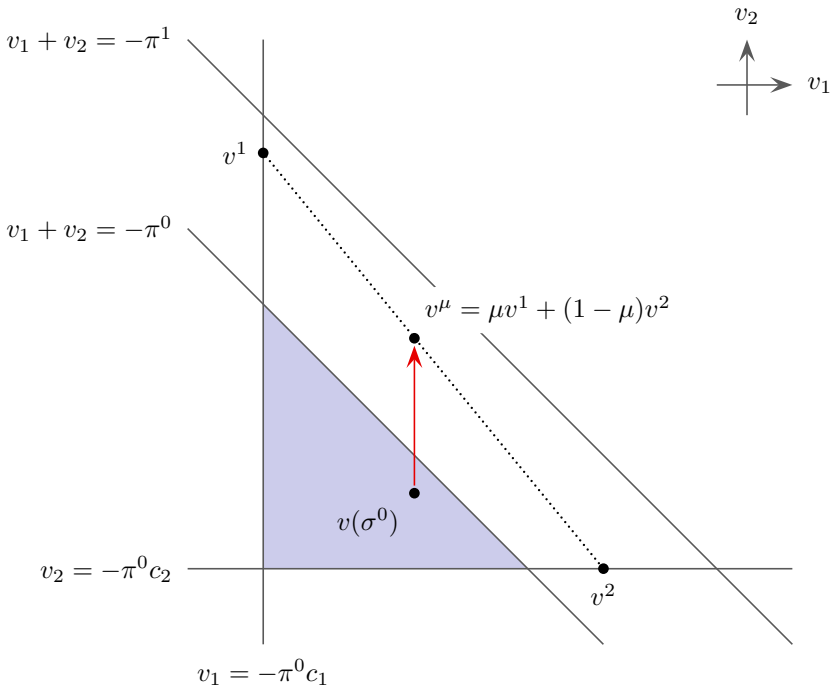

Figure 4: PPE without avoidance result in payoff profiles within the shaded area, all of which are dominated by convex combinations of v^1 and v^2.

Fix any $\lambda \in (0, \bar{\lambda})$, where

$$\bar{\lambda} := \min\left\{1,\ \frac{1-\delta}{\delta\Delta}(1-c_1)\right\} > 0.$$

Let $v^\lambda = (1-\lambda)u(0, c_1, 1-c_1) + \lambda v^*$. In particular, this implies that

$$v_1^\lambda = -\pi^0 c_1 + \lambda\Delta.$$

Since $\langle 0, c_1, 1-c_1\rangle$ is a PPE and $\mathcal{V}^*$ is convex (Proposition A.2), we know that $v^\lambda \in \mathcal{V}^*$ for all $\lambda \in (0,1)$.

Consider the action profile $s^\lambda = (a^\lambda, m_1^\lambda, m_2^\lambda)$, with $a^\lambda = 0$, $m_1^\lambda = c_1 + \varepsilon$, and $m_2^\lambda = 1 - c_1 - \varepsilon$, where

$$\varepsilon := \left(\frac{\delta}{1-\delta}\right)\lambda\Delta > 0.$$

Also, consider the profile of continuation values w^λ with $w_i^\lambda(0) = u_i(0, c_1, 1-c_1)$, and $w_i^\lambda(1) = v_i^\lambda$, for $i = 1, 2$. In words, (s^λ, w^λ) represents the following plan, see Figure 5. If there is no crisis, then the play transitions to $\langle 0, c_1, 1-c_1\rangle$ forever. If a crisis occurs, the play transitions to a mixture of $\langle 0, c_1, 1-c_1\rangle$ and the strategies that generate v^*. As we show below, the value of ε is carefully selected to guarantee that all the efficiency gains from avoidance go to agent 2, i.e., in order to have $g_1(s^\mu, w^\mu) = -\pi^0 c_1$.

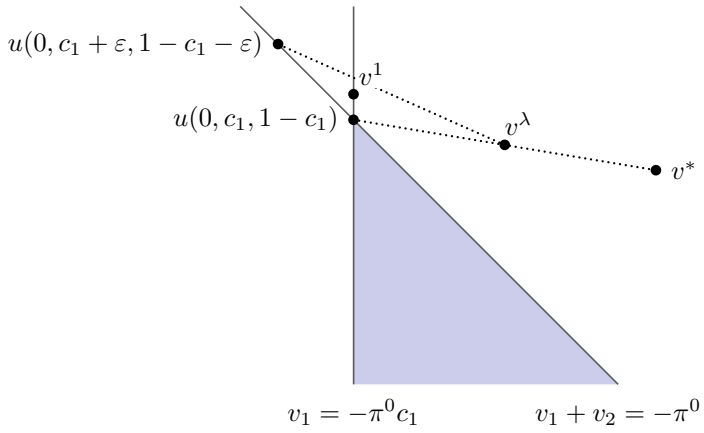

Figure 5: v^1 is a mixture of $u(0, c_1+\epsilon, 1-c_1-\epsilon)$ on the first period, $u(0, c_1, 1-c_1)$ from the second period onward if there is no crisis, and v^λ from the second period onward in case of a crisis.

The condition $\lambda < \bar{\lambda}$ guarantees that $c_1 + \varepsilon < 1$, so that $m_2^\lambda + m_1^\lambda = 1$, and m_2^λ is a static best response to m_1^λ. This implies that agent 2's mitigation constraint (M) for (s^λ, w^λ) is satisfied. Also, it implies that

$$(1-\delta)k_1(m_1^\lambda, m_2^\lambda) + \delta w_1^\lambda(1) = -(1-\delta)(c_1+\varepsilon) + \delta v_1^\lambda = -(1-\delta+\delta\pi^0)c_i.$$

This implies that agent 1's mitigation constraint and his avoidance constraint (A) are also satisfied, and that $g_1(s^\mu, w^\mu) = -\pi^0 c_1$. Therefore, (s^λ, w^λ) is admissible w.r.t. $\mathcal{V}^*$ and, by Proposition A.1, $g(s^\lambda, w^\lambda) \in \mathcal{V}^*$. Finally, note that

$$g_1(s^\mu, w^\mu) + g_2(s^\mu, w^\mu) = \pi^0\Big[(1-\delta)u_1(s^\lambda) + \delta v^\lambda\Big] - (1-\pi^0)\delta\pi^0 > -\pi^0.$$

Since $g_1(s^\mu, w^\mu) = -\pi^0 c_1$, this implies that $g_2(s^\mu, w^\mu) > -\pi^0(1-c_1)$, thus completing

the proof. ■

Finally, Lemma B.4 shows that it is not possible to have avoidance when the discount factor is very low, and it is possible to do so when it is very high. This, together with the monotonicity of the set of PPE payoffs with respect to the discount factor, implies the existence of the threshold δ^*—strictly between 0 and 1—separating a region where no avoidance is possible, from a region where avoidance and bailouts happen infinitely often in all efficient equilibria.

Lemma B.4. *There exist numbers $0 < \underline{\delta} < \bar{\delta} < 1$ such that if $\delta \leq \underline{\delta}$, then there is no avoidance in any PPE and if $\delta > \bar{\delta}$, then a PPE exists where the avoidance action is played with positive probability along the equilibrium path.*

Proof. Let $\bar{\delta}$ and $\underline{\delta}$ be the bounds for the discount factor given by

$$\bar{\delta} := \frac{\hat{d} - c_1}{\hat{d} - c_1 + \pi^0(c_1 + c_2 - 1)} \qquad \text{and} \qquad \underline{\delta} := \frac{\hat{d} - c_1}{\hat{d} - c_1 + \pi^0 c_1}.$$

Assumptions 1 and 2 require that $\hat{d} > c_1$, $c_1 + c_2 > 1$, and $c_2 < 1$. These conditions imply that $0 < \underline{\delta} < \bar{\delta} < 1$.

We begin by showing that, if $\delta > \bar{\delta}$, then a PPE exists where the avoidance action is played with positive probability. Let σ^0 be the public strategy profile described as follows. On the first period, $\sigma_1^0(h_1) = (1, c_1, 1 - c_1)$ for all $h_1 \in H_1$. If a crisis occurs on the first period, then $\sigma_t^0(h_t) = (0, c_1, 1 - c_1)$ for every subsequent history h_t. Otherwise, if there is no crisis on period one, then $\sigma_t^0(h_t) = (0, 1 - c_2, c_2)$ for every subsequent history h_t.

We will show that σ^0 is a PPE as long as $\delta \geq \bar{\delta}$. For $t > 1$, σ^0 consists of unconditional repetition of static Nash equilibria of the stage game. Hence, the continuation strategies are PPE, and we only need to check the incentive constraints from section A.1 for date $t = 1$. If there is a crisis in period 1, the agents' contributions are static mutual best responses and do not affect the continuation value. Therefore, the mitigation constraints (M) are satisfied. The avoidance constraint (A) can be written as

$$\begin{aligned} 0 &\leq \hat{d} + k_1(c_1, 1 - c_1) + \frac{\delta}{1 - \delta}\Big(u_1(0, 1 - c_2, c_2) - u_1(0, c_1, 1 - c_1)\Big) \\ &= \hat{d} - c_1 - \frac{\delta}{1 - \delta}\pi^0(c_1 + c_2 - 1). \end{aligned}$$

Rearranging terms, it is straightforward to show that this is equivalent to $\delta \geq \bar{\delta}$.

Now, we will show that if $\delta \leq \underline{\delta}$, then there cannot be any avoidance in any PPE. For that purpose, suppose that σ^* is a PPE with $\alpha_t^*(h_t) = 1$ for some history h_t. As in the proof of Proposition 4.1, the mitigation (M) and avoidance (A) constraints for the active agent after h_t can be written as

$$(1-\delta)k_1(m_1^*, m_2^*) + \delta w_1^*(1) \geq (1-\delta+\delta\pi^0)c_1$$
$$(1-\delta)k_1(m_1^*, m_2^*) + \delta w_1^*(1) \leq \delta w_1^*(0) - (1-\delta)\hat{d}.$$

Together, they imply that

$$(1-\delta)\hat{d} \leq (1-\delta+\delta\pi^0)c_1 + \delta w_1^*(0) \leq (1-\delta+\delta\pi^0)c_1,$$

where the last inequality follows from $w_1^*(0) \leq 0$. After some simple algebra, this condition is equivalent to $\delta \geq \underline{\delta}$. Hence, in order for a PPE with avoidance to exist, it cannot be the case that $\delta < \underline{\delta}$. ■

Now, we are in a position to prove Proposition 4.4.

Proof of Proposition 4.4. Let $D \subseteq (0,1)$ be the set of discount factors where it is possible to play avoidance, and let $\tilde{\delta} = \inf D$. Lemma B.4 implies that $D \neq \emptyset$. Lemma B.2 implies that for any $\delta \in D$, a PPE exists with total cost below π^0. Since the set of PPE payoffs is monotone with respect to the discount factor (Proposition A.2), the same is true for every $\delta' > \tilde{\delta}$. Since obtaining a cost below π^0 requires avoidance, this implies that $(\tilde{\delta}, 1) \subseteq D$. Lemma B.4 implies that $0 < \underline{\delta} \leq \tilde{\delta} \leq \bar{\delta} < 1$.

For the remainder of the proof, fix any $\delta > \tilde{\delta}$, and let σ^* be a PPE given δ. We are interested in the *first time* after which there is no more avoidance according to σ^*. For that purpose, let H^0 be the (possibly empty) set of histories h_t^0 after which there is no more avoidance, i.e., such that $a_{t'}^*(h_{t'}') = 0$ for every $h_{t'}'$ that follows h_t^0. Let $H' \subseteq H$ be the set of histories h_t' such that there is no $h_t^0 \in H^0$, which is a strict predecessor of h_t'.

For every history $h' \in H'$, let $\sigma^{h'}$ be the continuation strategies from h_t' onward. Since $\delta \in D$, we know by Lemma B.4 that the consequence of Lemma B.3 applies. Hence, $v^{h'} \in \mathcal{V}^*$ exists such that $v_1^{h'} = v_1(\sigma^{h'})$ and $v_2^{h'} > v_2(\sigma^{h'})$.

Now, let σ' be the strategy profile that mimics σ^* before it reaches any history

$h' \in H'$, and then it follows the strategies that support $v^{h'}$ instead of $\sigma^{h'}$. Since the continuation values after these histories for the active agent remain unchanged, and the continuation values for the passive agent increases, all the previous incentive constraints are still satisfied. Hence, σ' is a PPE that (weakly) Pareto dominates σ^*, and strictly Pareto dominates it if H' is reached with positive probability. Therefore, we can conclude that if σ^* is constrained efficient, then there is avoidance infinitely often almost surely. By Proposition 4.3, this implies that there are also bailouts infinitely often almost surely. ■

B.6. Proofs for the observable-actions model

Proof or Proposition 6.1. Proposition 6.1 is the analogous of Proposition 4.3 for the observable-actions case. Note that the proof of Proposition 4.3 did not make use of the fact that avoidance actions were unobservable. Hence, a completely analogous argument serves to prove Proposition 6.1. ■

Proof or Proposition 6.2. Let $\tilde{\delta}' \in (0,1)$ be the threshold given by

$$\tilde{\delta}' := \frac{d + \pi^1(1-c_2) - \pi^0 c_1}{d + \pi^1(1-c_2) - \pi^0 c_1 + \pi^1(d + \pi^0(c_1+c_2) - \pi^1)}. \tag{13}$$

Let σ^* be the grim-trigger strategy profile described in section 6.1. We will show that, if $\delta \geq \tilde{\delta}'$, then σ^* is a PPE of the repeated game with observable avoidance actions.

The single-deviation principle still applies, and hence it suffices to search for deviations at a single period. The strategies after a deviation are a repetition of a static Nash equilibrium of the stage game, and thus constitute a PPE. Along the equilibrium path, it is optimal for the active agent to choose the avoidance action if and only if:

$$-(d + \pi^1 m_1^*) \geq -\pi^0 c_1.$$

The mitigation incentive constraint for agent 1 is:

$$-(1-\delta)m_1^* - \delta(d + \pi^1 m_1^*) \geq -(1-\delta)c_1 - \delta\pi^0 c_1,$$

and for agent 2 it is:

$$-(1-\delta)m_2^* - \delta\pi^1 m_2^* \geq -(1-\delta)c_2 - \delta\pi^0 c_2,$$

where $m_2^* = 1 - m_1^*$. The mitigation constraint for agent 2 is satisfied if and only if:

$$m_2^* \leq \left(\frac{1-\delta+\delta\pi^0}{1-\delta+\delta\pi^1}\right) c_2.$$

Hence, we can choose m_2^* to satisfy this condition with equality, and incentives for player 2 are automatically satisfied. Then, using the fact that $m_1^* + m_2^* = 1$, we have that:

$$m_1^* = 1 - \left(\frac{1-\delta+\delta\pi^0}{1-\delta+\delta\pi^1}\right) c_2 < 1 - c_2 < c_1.$$

This implies that $c_1 - m_1^* > 0$, and thus, after some simple algebra, the mitigation constraint for 1 is satisfied whenever the avoidance constraint is satisfied. This implies that the proposed strategy profile is a PPE if and only if:

$$d + \pi^1 - \pi^1 \left(\frac{1-\delta+\delta\pi^0}{1-\delta+\delta\pi^1}\right) c_2 \leq \pi^0 c_1,$$

which, after some algebra, is equivalent to $\delta \geq \tilde{\delta}'$. ■

B.7. Proofs for the monetary-transfers model

Proof or Proposition 6.3. Proposition 6.3 is the analogous of Proposition 4.3 for the monetary-transfers model. The argument of the proof is analogous. Consider a strategy profile with no bailouts or monetary compensation, i.e., such that for almost every history (h_t), $\beta_{t0}(h_t) = 0$ and $\beta_{t1}(h_t) - \mu_{1t}(h_t) \leq -c_1$. This implies that the per-period utility for the active agent is less or equal than $-d - \pi^1 c_1$ for histories with $\alpha_t(h_t) = 1$, and less or equal than $-\pi^0 c_1$ for histories with $\alpha_t(h_t) = 0$. Assumption 2 implies that $-d - \pi^1 c_1 < -\pi^0 c_1$, and thus $v_1(\sigma) \leq -\pi^0 c_1$, with strict inequality whenever $\alpha_t(h_t) = 1$ for some history set of histories $\{h_t\}$, which is reached with positive probability along the equilibrium path. Individual rationality requires $v_1(\sigma) \geq -\pi^0 c_1$. Hence, a strategy profile with no bailouts can satisfy individual rationality only if

$\alpha_t(h_t) = 0$ almost surely along the equilibrium path. ■

Proof of Proposition 6.4. Let σ^* be the grim-trigger strategy profile described in section 6.2, with the constants b^* and m_1^* taking the values

$$m_1^* := c_1 + \delta(1-\pi^0)(\hat{d} - c_1) \qquad \text{and} \qquad b^* = (1 - \delta + \delta\pi^0)(\hat{d} - c_1). \tag{14}$$

Assumption 2 guarantees that $m_1^* \in (0,1)$ and $b^* > 0$. We will show that, if $\delta \geq \tilde{\delta}''$, then σ^* is a PPE of the repeated game with observable avoidance actions. The single-deviation principle still applies. The strategies after a detectable deviation are a repetition of a static Nash equilibrium of the stage game, and thus constitute a PPE. Hence, we only need to verify that there are no profitable deviations along the equilibrium path.

Since the avoidance action is not observable and s^* is stationary, if it were profitable for the active agent to deviate *at a single period* by not taking the avoidance action on that period, then it would also be profitable to deviate by not taking the avoidance action *in any period*. Hence, in order to show that it is optimal for the active agent to take the avoidance action along the equilibrium path, it suffices to show that his equilibrium average expected discounted utility is weakly greater than the average expected discounted utility he would get by taking $a = 0$ on every period. We denote the value of this deviation v_1'. Note that $b^* = \hat{d} - m_1^*$. As we show below, this value was chosen specifically so that $v_1' = v_1(\sigma^*)$. Since σ^* is stationary along the equilibrium path, each agent's average expected discounted utility equals their per-period utility. In particular, for the active agent we have that

$$\begin{aligned} v_1(\sigma^*) &= -d - \pi^1 m_1^* + (1-\pi^1)b^* = -d - \pi^1 m_1^* + (1-\pi^1)(\hat{d} - m_1^*) \\ &= -m_1^* + d\left(\frac{1-\pi^1}{\pi^0 - \pi^1} - 1\right) = -m_1^* + (1-\pi^0)\hat{d}. \end{aligned} \tag{15}$$

By a similar argument, we have that

$$v_1' = -\pi^0 m_1^* + (1-\pi^0)b^* = (1-\pi^0)\hat{d} - m_1^*,$$

which implies that $v_1' = v_1(\sigma^*)$.

In turn, using an analogous argument for the benchmark model (see section A.1), the mitigation constraint for the active agent on the monetary-transfer model can be

written as

$$\begin{aligned} & -(1-\delta)m_1^* + \delta v_1(\sigma^*) \geq -(1-\delta+\delta\pi^0)c_1 \\ \Leftrightarrow \quad & -(1-\delta)m_1^* + \delta(-m_1^* + (1-\pi^0)\hat{d}) \geq -(1-\delta+\delta\pi^0)c_1 \\ \Leftrightarrow \quad & c_1 + \delta(1-\pi^0)(\hat{d}-c_1) \geq m_1^*, \end{aligned}$$

where the first implication follows from (15), and the second implication is obtained by rearranging terms. From (14), it follows that the mitigation constraint for the active agent is satisfied with equality.

The values of m_1^* and b^* were specifically chosen so that the incentive constraints for the active agent are satisfied with equality. Now, it remains to show that, when δ is close enough to 1, this leaves enough slack for the passive agent to be able to contribute $m_2^* = 1 - m_1^*$ in case of a crisis and to transfer b^* to the passive agent each time there is no crisis. The mitigation constraint for the passive agent is

$$-(1-\delta)m_2^* + \delta v_2(\sigma^*) \geq -(1-\delta)c_2 - \delta\pi^0 c_2,$$

and his constraint for transfers in case there is no crisis is

$$-(1-\delta)b^* + \delta v_2(\sigma^*) \geq -\delta\pi^0 c_2.$$

Below, we will show that $\lim_{\delta\uparrow 1} v_2(\sigma^*) > -\pi^0 c_2$. This implies that, in the limit when the discount factor approaches 1, both constraints are satisfied with strict inequality. By continuity, this implies that there exists some $\tilde{\delta}''$ such that σ^* is a PPE of the monetary-transfers model as long as $\delta \geq \tilde{\delta}''$.

By a similar argument as before, the average expected equilibrium value for the passive agent is

$$\begin{aligned} v_2(\sigma^*) &= -\pi^1(1-m_1^*) - (1-\pi^1)b^* = -\pi^1 + m_1^* - (1-\pi^1)\hat{d} \\ &= -\pi^1 - (1-\pi^1)\hat{d} + c_1 + \delta(1-\pi^0)(\hat{d}-c_1). \end{aligned}$$

Therefore, we have that

$$\lim_{\delta\uparrow 1} v_2(\sigma^*) = -\pi^1 - (1-\pi^1)\hat{d} + c_1 + (1-\pi^0)(\hat{d}-c_1) = \pi^0 c_1 - \pi^1 - d. \tag{16}$$

Finally, Assumptions 1 and 2 imply that

$$\begin{aligned} & c_1 + c_2 > 1 \qquad \wedge \qquad d < \pi^0 - \pi^1 \\ \Rightarrow \quad & \pi^0(c_1 + c_2) - \pi^1 > \pi^0 - \pi^1 > d \\ \Rightarrow \quad & \pi^0 c_1 - \pi^1 - d > -\pi^0 c_2. \end{aligned}$$

Hence, by (16), we have $\lim_{\delta \uparrow 1} v_2(\sigma^*) > -\pi^0 c_2$, thus completing the proof. ■

References

Abreu, Dilip, David Pearce, and Ennio Stacchetti. 1990. "Toward a theory of discounted repeated games with imperfect monitoring." *Econometrica* 58 (5):1041–1063.

Chari, VV and Patrick J Kehoe. 2015. "Bailouts, time inconsistency, and optimal regulation: A macroeconomic view." *American Economic Review* .

Cronshaw, Mark B. and David G. Luenberger. 1994. "Strongly symmetric subgame perfect equilibria in infinitely repeated games with perfect monitoring and discounting." *Games and Economic Behavior* 6 (2):220–237.

Diamond, Douglas W and Philip H Dybvig. 1983. "Bank runs, deposit insurance, and liquidity." *Journal of Political Economy* :401–419.

Ennis, Huberto M and Todd Keister. 2009. "Bank runs and institutions: The perils of intervention." *American Economic Review* 99 (4):1588–1607.

Farhi, Emmanuel and Jean Tirole. 2012. "Collective moral hazard, maturity mismatch, and systemic bailouts." *American Economic Review* 102 (1):60–93.

Green, Edward J. 2010. "Bailouts." *FRB Richmond Economic Quarterly* 96 (1):11–32.

Green, Edward J. and Robert H. Porter. 1984. "Noncooperative collusion under imperfect price information." *Econometrica* 52 (1):87–100.

Keister, Todd. 2016. "Bailouts and financial fragility." *Review of Economic Studies* 83 (2):704–736.

Mailath, George J. and Larry Samuelson. 2006. *Repeated games and reputations: Long-run relationships*. Oxford University Press.

Radner, Roy. 1985. "Repeated principal-agent games with discounting." *Econometrica* 53 (5):1173–1198.

Rubinstein, Ariel and Menahem E. Yaari. 1983. "Repeated insurance contracts and moral hazard." *Journal of Economic Theory* 30 (1):74–97.

Schneider, Martin and Aaron Tornell. 2004. "Balance sheet effects, bailout guarantees and financial crises." *Review of Economic Studies* 71 (3):883–913.