Online Appendix for
"Institution Building without Commitment"

Marco Bassetto   Zhen Huo   José-Víctor Ríos-Rull

A Proofs of Section III.A

A.1 Proposition A.1

Proposition A.1. *Under Assumptions 1 and 2, there exists a subgame-perfect equilibrium of the game that satisfies Requirements 1 and 2.*

Proof. Assumption 2 follows Kocherlakota (1996), who uses it in Proposition 4 to prove that a reconsideration-proof equilibrium exists for the game whose period-t payoff is $V(a_t, a_{t+1}, a_{t+2}, ...)$. The strategies of such a game represent a subgame-perfect equilibrium of our game with a state variable: weak separability implies that the state does not affect the preference ordering of each player over the sequence of future actions. Moreover, these strategies satisfy Requirements 1 and 2 by the definition of a reconsideration-proof equilibrium. □

A.2 Proof of Proposition 1

Proof. Let $(a_0^E, a_1^E, \ldots)$ be the outcome of a reconsideration-proof equilibrium for the game whose period-t payoff is $V(a_t, a_{t+1}, a_{t+2}, ...)$, and let $\bar{V}$ be its associated value. This means that, for any period t and any actions $a \in A$, there exists a continuation sequence $(a_{t+1}^{\bar{E}}, a_{t+2}^{\bar{E}}, \ldots)$ which is also the outcome of a reconsideration-proof equilibrium and is such that

$$V(a_t^E, a_{t+1}^E, a_{t+2}^E, \ldots) \geq V(a, a_{t+1}^{\bar{E}}, a_{t+2}^{\bar{E}}, \ldots). \tag{1}$$

We then have

$$V(a, a_{t+1}^{\bar{E}}, a_{t+2}^{\bar{E}}, \ldots) = \widetilde{V}(a, \widehat{V}(a_{t+1}^{\bar{E}}, a_{t+2}^{\bar{E}}, \ldots)).$$

Acknowledging that the sequence $(a_{t+1}^{\bar{E}}, a_{t+2}^{\bar{E}}, \ldots)$ is potentially a function of the deviation a (as well as of time t, which we can hold fixed), define

$$\underline{V} := \inf_{a \in A} \widehat{V}(a_{t+1}^{\bar{E}}, a_{t+2}^{\bar{E}}, \ldots). \tag{2}$$

By the compactness of A, Tychonoff's theorem, and continuity of $\widehat{V}$, we can find a sequence of actions $a_0^*, a_1^*, ...$ that attains the infimum in equation (2) above. Exploiting Assumption 3, this sequence ensures subgame perfection and satisfies the no-restarting condition (Requirement 3):

$$V(a_0^*, a_1^*, a_2^*, \ldots) \geq V(a, a_0^*, a_1^*, \ldots).$$

This path attains the value $\bar{V}$, so that it continues to satisfy the optimality condition of Requirement 2. Hence, playing $(a_0^*, a_1^*, ...)$ followed by a restart after any deviation is an equilibrium that satisfies Requirements 1, 2, and 3, and therefore $(a_0^*, a_1^*, ...)$ is an organizational equilibrium. □

B Proofs and Details in Section III.B

The players in the game are nature, plus an infinity of players 0,1,... indexed by the time at which they act. Nature moves first, choosing a time $\hat{t}$, from which record keeping is possible. We assume that this distribution has full support on $\mathbb{N}$.[1]

Players take two actions:

- Player t chooses $a_t \in A$.
- In addition, a player may choose a record-keeping action $\rho_t \in \{S, C, H\}$, where S stands for starting record keeping, C stands for continuing record keeping, and H for hiding past records. Whether these actions are available at time t depends on the past in a way that we will soon make explicit.

We are now ready to define histories and information. The first history is $\emptyset$, at which stage nature moves. In all periods $t < \hat{t}$, players only choose the action a_t. While the history of play is $(\hat{t}, a_0, ..., a_{t-1})$, their only information is that $t < \hat{t}$, and the current level of the state k_t: they do not observe any of the past players' actions, and they only know that record keeping is not yet possible. In period $t = \hat{t}$, the history of play is also $(\hat{t}, a_0, ..., a_{t-1})$. Player $t = \hat{t}$ observes k_t, does not observe any of the actions taken by the past players, but it knows that $t \geq \hat{t}$ and that either $t = \hat{t}$ or $\rho_{t-1} = H$: that is, it knows that it is either the first player with the opportunity to set up record keeping, or the opportunity was available in the past, but player $t-1$ chose not to adopt it and to hide the previous history. Player $t = \hat{t}$ is called to choose an action $\rho_t \in \{S, H\}$ as well as a_t. In period $\hat{t}+1$ and all subsequent periods, the history of play is $(\hat{t}, a_0, ..., a_{\hat{t}-1}, \rho_{\hat{t}}, a_{\hat{t}}, \rho_{\hat{t}+1}, a_{\hat{t}+1}, ..., \rho_{t-1}, a_{t-1})$. In each of these periods, if $\rho_{t-1} = H$, then player t only knows that record keeping is possible and the level of k_t; it does not know whether $t = \hat{t}$ or $\rho_{t-1} = H$. In this case, player t has the same options as player $\hat{t}$. Otherwise, let $\tilde{t}$ be the last time action S was taken; player t then knows that $\hat{t} \leq \tilde{t}$ and it knows $(\rho_{\tilde{t}}, a_{\tilde{t}}, \rho_{\tilde{t}+1}, a_{\tilde{t}+1}, ..., \rho_{t-1}, a_{t-1})$ (in addition to k_t). Player t has 3 options for ρ_t: first, it can choose $\rho_t = H$, in which the next player will start again with no record of the past; second, it can choose $\rho_t = C$, that is, to continue record keeping: in this case, player $t+1$ will know $(\rho_{\tilde{t}}, a_{\tilde{t}}, \rho_{\tilde{t}+1}, a_{\tilde{t}+1}, ..., \rho_t, a_t)$. Finally, it can restart the history ($\rho_t = S$), disavowing the past, but recording its own actions, in which case player $t+1$ will only observe (ρ_t, a_t). In all cases player $t+1$ will observe k_{t+1}.

A strategy σ_t for player t is a mapping from the set of time-t histories, H^t, to the set of actions A and (when available) record-keeping choices $\rho_t \in \{H, S, C\}$, that is measurable with respect to the information available at time t. As before, a strategy profile σ is a sequence of strategies, one for each player. It is useful to distinguish between the two choices made by agents: accordingly, let $\sigma_{a\,t}$ be the component of $\sigma_t(h^t)$ that contains the prescribed action $a \in A$ after history h^t, and $\sigma_{\rho\,t}(h^t)$ be the prescribed choice of record keeping. Analogously, we define $\sigma_a := \{\sigma_{a\,t}\}_{t=0}^{\infty}$ and $\sigma_\rho := \{\sigma_{\rho\,t}\}_{t=0}^{\infty}$.

We restrict attention to equilibria that satisfy Requirement 1: that is, they involve strategies that are inde-

[1]It would be equivalent to assume that nature moves in each period up to $\hat{t}$, as long as the conditional hazard rate of the start of record keeping is the same. This is because nature's choice is not fully observed by the agents anyway.

pendent of k_t.

We define a *full-disclosure* equilibrium to be an equilibrium in which $\sigma_{\rho\, t}(h^t) = S$ for all histories for which $t \geq \hat{t}$ and no previous record of play is known, and $\sigma_{\rho\, t}(h^t) = C$ for all histories for which player t observes a record of past actions. In a full-disclosure equilibrium, players introduce record keeping as soon as possible and never erase any of the record available, independent of the actions of past players.

The proof of Proposition 2 relies on a sequence of lemmata:

Lemma 1. *Let σ be a sequential equilibrium that satisfy Requirement 1 in the game defined above. Then:*

1. *There exists a full-disclosure sequential equilibrium $\tilde{\sigma}$ that also satisfies that satisfy Requirement 1 and such that the same actions $\{a_t\}_{t=0}^{\infty}$ are taken on the equilibrium path under σ and $\tilde{\sigma}$.*
2. *If σ satisfies Requirement 2 from period $\hat{t}$ (whatever $\hat{t}$ turns out to be), then $\tilde{\sigma}$ can be chosen to also satisfy the same requirement.*

Proof. Our proof only looks at pure-strategy equilibria. It could be extended to mixed-strategy equilibria, in which players randomize over their choice of record keeping, using the same logic presented here, as long as a public randomization device is present that allows coordination across players. We omit the case of mixed-strategy equilibria for brevity.

1. Assumption 1 implies that, if future players do not condition their choices on the state k (but potentially condition their choices on all their remaining information in any arbitrary way), the optimal choice for a current player is independent of the current state. In looking at equilibria that satisfy Requirement 1, we can therefore leave the state k in the background and focus only on the history of actions, disclosures, and the time at which record keeping becomes available.

   Let $\sigma = \{\sigma_t\}_{t=0}^{\infty}$ be the strategy profile of the sequential equilibrium that contains the equilibrium action path $\{a_t\}_{t=0}^{\infty}$.

   We need to construct an alternative strategy profile $\tilde{\sigma}$ that contains the same equilibrium action path, but involves full disclosure. We will do so by creating a suitable mapping from the set of histories to itself, and setting $\tilde{\sigma}_{a\, t}(h^t) = \sigma_{a\, t}(\eta(h^t))$. η is constructed recursively as follows:

   - For $t \leq \hat{t}$, $\eta(h^t) = h^t$.
   - For $t > \hat{t}$ and histories in which $\rho_{t-1} = H$, $\eta(h^t) = h^t$.
   - For $t > \hat{t} + 1$ and histories in which $\rho_{t-1} = S$ and $\sigma_{\rho,t}(h^{t-1}) = S$ or $\sigma_{\rho,t}(h^{t-1}) = C$, $\eta(h^t) = h^t$.
   - For $t > \hat{t}+1$ and histories in which $\rho_{t-1} = S$ and $\sigma_{\rho,t}(h^{t-1}) = H$, $\eta(h^t) = (h^{t-1}, H, a_{t-1,h^t})$, where a_{t-1,h^t} is the action taken in period $t-1$ according to the history h^t.
   - For $t > \hat{t}$ and histories in which $\rho_{t-1} = C$, we define η recursively as $\eta(h^t) = (\eta(h^{t-1}), \sigma_{\rho,t}(h^{t-1}), a_{t-1,h^t})$.

Furthermore, whenever $t \geq \hat{t}$, $\tilde{\sigma}_{\rho\, t} = S$ if no record keeping is currently in place, and $\tilde{\sigma}_{\rho\, t} = C$ otherwise, in line with the definition of a full-disclosure equilibrium.

In words, $\tilde{\sigma}$ is constructed from σ by assuming that agents take the same actions under the two strategy profiles whenever they do not observe the past. When past actions are observed from $\bar{s}$ on, the strategy profile $\tilde{\sigma}$ prescribes that the agents take the same actions they would have taken under σ when faced with a history that has same choices for $(a_0, ..., a_{t-1})$, but in which past players from $\bar{s}$ on chose to hide, start, or continue record-keeping according to the equilibrium profile σ. At the same time, $\tilde{\sigma}$ always prescribes full disclosure. Next, we verify that $\tilde{\sigma}$ is a measurable strategy with respect to the information sets available to the players at each point t. The choice of ρ only depends on whether record keeping is possible and whether it is inherited from the past, which is observable to an agent at the time it makes its choice. Furthermore, by construction, the mapping η is such that the prescribed action $\tilde{\sigma}_{a\, t}(h^t)$ is the same for all histories that share the same observable record.[2]

Next, we verify that $\tilde{\sigma}$ represents a sequential equilibrium. A player's payoff only depends on the current and future actions $a_t \in A$, and only indirectly on record keeping choices.

- In any period $t < \hat{t}$, the current choice of a_t by player t is not known to future players and therefore it has no impact on any future action. Furthermore, the two strategies σ and $\tilde{\sigma}$ imply the same sequence of future actions $(a_{t+1}, a_{t+2}, ...)$ along the equilibrium path.[3] The optimality of $\tilde{\sigma}_t$ then follows directly from that of σ_t.
- Consider next periods $t \geq \hat{t}$ and histories h^t such that no record is available to player t. For such histories, $\eta(h^t) = h^t$. There are two possibilities. First, suppose that $\sigma_\rho(h^t) = S$. Then, no matter what choice of (ρ_t, a_t) player t takes, the equilibrium implies that future players will take the same actions $\{a_s\}_{s=t+1}^{\infty}$ under profiles σ and $\tilde{\sigma}$. Hence, $\tilde{\sigma}_t(h^t) = \sigma_t(h^t)$ is an optimal choice. Suppose instead that $\sigma_{\rho\, t}(h^t) = H$, that is, according to the equilibrium profile σ, player t should hide its action. In this case, η is such that player t gets the same payoff whether it chooses $\rho_t = S$ or $\rho_t = H$, since $\eta(h^t, H, a_t) = \eta(h^t, S, a_t)$: player t is indifferent between starting record keeping or not, because in either case future players will ignore its play and behave as if no record had been taken in t. Starting record keeping is thus weakly optimal, and taking the same action that would have been taken under the profile σ is optimal as well.[4]
- Consider histories h^t in which a record is present. The reasoning is similar. If $\sigma_{\rho,t}(\eta(h^t)) = C$, then, no matter what choice of (ρ_t, a_t) player t takes, the equilibrium implies that future players will take the same actions $\{a_s\}_{s=t+1}^{\infty}$ under profiles $\tilde{\sigma}$ and history h^t as they would under σ and history $\eta(h^t)$. Hence, if $\sigma_t(\eta(h^t))$ is optimal (taking as given that σ will be followed in the future), then $\tilde{\sigma}_t(h^t)$ is also optimal, if future players play according to $\tilde{\sigma}$. If $\sigma_{\rho\, t}(\eta(h^t)) = H$, then under $\tilde{\sigma}$ future players will ignore past actions whether player t chooses $\rho_t = H$ or $\rho_t = C$, and their future actions will follow the course dictated by $\sigma|_{(h^t,H,a_t)}$. By the measurability restriction, $\sigma|_{(h^t,H,a_t)} = \sigma|_{(\eta(h^t),H,a_t)}$. Hence, player t is indifferent between playing C or H. If player t

[2] This assumes that the property is true for σ, which must be the case for σ to be a valid strategy profile and therefore a valid equilibrium, provided that σ does not condition on k_t, which is guaranteed by Requirement 1.

[3] Notice that future actions are in general uncertain and depend on the realization of $\hat{t}$, but their stochastic process is identical in the two equilibria.

[4] Since future players will ignore the action a_t, player t will maximize its payoff assuming that its action does not affect the future, as if no record were taken, just as it would under the strategy σ, which prescribes hiding the record.

chooses to restart the record, then the future players' actions will evolve according to $\sigma|_{(h^t,S,a_t)}$. Measurability implies again that $\sigma|_{(h^t,S,a_t)} = \sigma|_{(\eta(h^t),S,a_t)}$. Since σ is an equilibrium profile, playing $\sigma_t(\eta(h^t))$ (which in this case involves hiding the record from future players) is weakly better than playing S along with any of the possible actions, under the assumption that future players will follow the same profile σ. It follows that the consequences of playing H vs. S and any action a_t in period t under history h^t when future players will follow $\tilde{\sigma}$ are the same as those of playing the corresponding actions under history $\eta(h^t)$ when future players will follow σ. Hence, if $\sigma(\eta(h^t)) = H$, playing S is a (weakly) dominated choice. In sum, in this case player t is indifferent between H and C, and it weakly prefers either to S, which ensures that it is optimal for its to play C. Furthermore, choosing $a_t = \tilde{\sigma}_{a\,t}(h^t) = \sigma_{a\,t}(\eta(h^t))$ is optimal because it involves a static optimization taking as given the future choices (that will be independent of the current a_t and will be the same under h^t and $\tilde{\sigma}$ as they are under $\eta(h^t)$ and σ). The last case to consider is one in which $\sigma_{\rho\,t}(\eta(h^t)) = S$; this case is similar to the previous one. Specifically, the measurability restriction implies $\sigma|_{(h^t,S,a_t)} = \sigma|_{(\eta(h^t),S,a_t)}$. Furthermore, if player t chooses $\rho_t = C$, $\tilde{\sigma}$ is such that future players will choose the same sequence of actions whether player t chooses $\rho_t = S$ or $\rho_t = C$: these actions will only depend on a_t, which is the only element of the record that is passed to future players according to the strategy σ_t. If player t chooses $\rho_t = H$, the future equilibrium path unfolds according to $\tilde{\sigma}|_{(h^t,H,a_t)} = \sigma|_{h^t,H,a_t} = \sigma|_{(\eta(h^t),H,a_t)}$, where the last equality follows the usual measurability restriction. If $\sigma_t(\eta(h^t)) = S$, then playing $\rho_t = S$ is weakly better than playing $\rho_t = C$ at $\eta(h^t)$ if σ will be followed in the future; this then implies that S (and the best action a_t conditional on S) is weakly better than C (and the best a_t conditional on C) at history h^t if $\tilde{\sigma}$ will be played in the future. This establishes that, under $\tilde{\sigma}$, playing C yields the same payoff as playing S, and a weakly better payoff than playing H. So, playing C is optimal. Finally, the usual equivalence of future consequences implies that playing $a_t = \tilde{\sigma}_{a\,t}(h^t) = \sigma_{a\,t}(\eta(h^t))$ is optimal.

2. Note that $\tilde{\sigma}$ is constructed so that the actions on the equilibrium path starting from any history h^t (whether the history itself is on or off equilibrium) are the same as the actions on the equilibrium path starting from $\eta(h^t)$ when σ is played. The mapping η is such that histories with $t \geq \hat{t}$ are mapped into histories with $t \geq \hat{t}$. If V is symmetric, then it achieves the same action payoff V following any history that has $t \geq \hat{t}$; as a consequence, the same property is inherited by $\tilde{\sigma}$. This implies that the set of values attainable by sequential equilibria satisfying Requirement 1 from period $\hat{t}$ is the same as the set of values attainable by full-disclosure sequential equilibria satisfying Requirement 1 from the same period; the maxima of the two sets will thus coincide, completing the proof.[5]

□

Lemma 2. *Let $\tilde{\sigma}$ be a full-disclosure state-independent sequential equilibrium for the game in which history can be hidden. Then:*

1. $\tilde{\sigma}_a|_{h^{\hat{t}}} \equiv \sigma$ is a subgame-perfect equilibrium for the game where record-keeping starts at time 0, and it

[5]The inability to keep records for periods before $\hat{t}$ will in general imply that the payoff in previous periods is lower.

also satisfies state independence (Requirement 1);[6]

2. *If $\tilde{\sigma}$ is symmetric from period $\hat{t}$ on, then $\tilde{\sigma}_a|_{h^{\hat{t}}}$ is also symmetric.*

Proof. 1. In the game in which history can be hidden, in period $\hat{t}$, player $\hat{t}$ starts with no information about the past, just as in period 0 of the game where record-keeping starts at time 0. Furthermore, $\tilde{\sigma}$ is such that records will be kept from $\hat{t}$ on. Take as given the choice of ρ_t dictated by $\tilde{\sigma}$, and focus on the choice of a_t. In order for $\tilde{\sigma}$ to represent a sequential equilibrium, at any time $t \geq \hat{t}$ and after any history h^t it must be the case that $\sigma_{a\,t}(h^t)$ (along with starting record keeping if no record is present or continuing it otherwise) is optimal, conditional on the fact that future players will continue to play $\tilde{\sigma}$. Let $h^t_{a,\hat{t}}$ represent the subcomponent of history h^t that captures the history of actions $(a_{\hat{t}}, a_{\hat{t}+1}, ...a_t)$. Since $\tilde{\sigma}$ implies that future players will behave in such a way that the entire history of play from $\hat{t}$ is known, it then follows that $\sigma_{a\,t}(h^t)$ must be optimal in the game where record keeping starts in period 0 after history $h^t_{a,\hat{t}}$, assuming that future players will play according to the strategy profile $\tilde{\sigma}_a|_{h^{\hat{t}}}$.

2. Symmetry implies that the action payoff V on the equilibrium path conditional on attaining any history h^t with $t \geq \hat{t}$ is the same. This property is inherited by $\tilde{\sigma}_a|h^{\hat{t}}$ in any subgame following a history $h^t_{a,\hat{t}}$, since the action paths coincide going forward.

□

Lemma 3. *Let σ be a symmetric state-independent subgame-perfect equilibrium of the game where record-keeping starts in period 0. Then, if and only if σ satisfies Requirement 3 as well, there exists a state-independent full-disclosure sequential equilibrium $\tilde{\sigma}$ of the game where history can be hidden, which is symmetric from period $\hat{t}$ and is such that $\tilde{\sigma}_a|_{h^{\hat{t}}} \equiv \sigma$.*

Proof. Assume first that σ satisfies Requirement 3. The condition $\tilde{\sigma}_a|_{h^{\hat{t}}} \equiv \sigma$ fully characterizes $\tilde{\sigma}_a$ from period $\hat{t}$ on. To see this, let h^t be an arbitrary history in which $t > \hat{t} + s$ and player t observes $(a_{t-s}, a_{t+1-s}, ..., a_{t-1})$; this implies that either player $t-s-1$ chose to hide records, or player $t-s$ chose to restart them, while all subsequent players up to t chose to continue record keeping. This history is in the same information set as a history with the same sequence of actions $(a_{t-s}, a_{t+1-s}, ..., a_{t-1})$ in which $\hat{t} = t - s$ and players adopted full disclosure; actions for such history are determined by $\tilde{\sigma}_a|_{h^{\hat{t}}} \equiv \sigma$. This observation also implies that, following any history, the sequence of actions a that are predicted to happen along a continuation equilibrium according to $\tilde{\sigma}$ is the same as those in a corresponding history in the game where record-keeping starts in period 0 under σ. If all histories under σ are followed by the same equilibrium action payoff $\bar{V}$, then the same value carries over to $\tilde{\sigma}$. To verify that $\tilde{\sigma}$ is indeed optimal after any history $h^t, t \geq \hat{t}$, we denote $h^t_{a\,s} = (a_s, ..., a_t)$ to be the record available to player t after history h^t and proceed as follows:[7]

[6] Note that, without further assumptions, $\tilde{\sigma}_a|_{h^{\hat{t}}}$ may depend on the precise realization of $\hat{t}$. The property still holds: in this case, each possible continuation strategy $\tilde{\sigma}_a|_{h^{\hat{t}}}$ is a subgame-perfect equilibrium of the game where record-keeping starts at time 0.

[7] Along the equilibrium path, the record available should start from period $\hat{t}$, but we need to verify optimality even for histories that are not on the equilibrium path.

- Player t does not have an incentive to choose $\rho_t = C$ and any action $a \neq \sigma_t(h^t_{a,\hat{t}})$. Assuming that future players will follow $\tilde{\sigma}$, the consequences of such a choice would be the same as those of choosing $a \neq \sigma_t(h^t_{a,\hat{t}})$ after history $h^t_{a,\hat{t}}$ in the game where record-keeping starts in period 0 when future players follow σ; since σ represents an equilibrium, choosing $a \neq \sigma_t(h^t_{a,\hat{t}})$ is weakly worse.

- Player t does not have an incentive to choose $\rho_t = S$ and any action $a \in A$. Following such a choice, player $t+1$ will behave as if $\hat{t} = t$, and future actions will unfold according to the strategy profile σ. Requirement 3 implies that, whatever action player t chooses, it would be (weakly) better off playing $\rho_t = S$ and $a = \sigma(\emptyset)$, that is, choosing to restart record keeping and playing the first action of the strategy profile of the game where record-keeping starts in period 0. This latter choice gives an action payoff of $\bar{V}$, which is the same as that obtained by continuing record keeping and following $\tilde{\sigma}$.

- Player t does not have an incentive to choose $\rho_t = H$ and any action $a \in A$. Following such a choice, player $t+1$ and subsequent players will follow the strategy σ as if the game in which record-keeping starts in period 0 took place from that point on. Requirement 3 implies that, faced with this prospect, player t does not have any action that can guarantee a payoff higher than $\bar{V}$ for herself.

To finish establishing the “if" part of the Lemma, the last step is to construct the strategy profile $\tilde{\sigma}$ in periods $t < \hat{t}$. In these periods, the actions taken by player t will not be observed by future players; as long as $\tilde{\sigma}$ is independent of the state, the actions of the current player will thus have no consequences on the actions taken by future players. We thus need to prove existence of a sequence of actions $(\tilde{a}_0, \tilde{a}_1, ...)$ that will be taken by players in period t if $t < \hat{t}$, and that are optimal given that the same sequence will be continued up to the unknown time $\hat{t}$ and given that starting in period $\hat{t}$ actions will unfold according to the equilibrium path dictated by σ. Given σ, consider a correspondence $M : A^\infty \rightrightarrows A^\infty$ that associates to a sequence $(a_0, a_1, ...)$ all the sequences such that player t is choosing optimally given that $(a_0, a_1, ...)$ will be followed up to period $\hat{t}$ and σ will be followed from period $\hat{t}$ on. By Assumptions 2 and 4 and the theorem of the maximum, M is nonempty, compact- and convex-valued, upper hemicontinuous, and independent of the state. By Kakutani's fixed-point theorem, M has a fixed point, which can be used as our desired sequence $(\tilde{a}_0, \tilde{a}_1, ...)$.

Conversely, suppose that σ *does not* satisfy Requirement 3. We know from the previous part of the proof that player $t \geq \hat{t}$ can attain the action payoff $\bar{V}$ by continuing record keeping and following the strategy $\tilde{\sigma}$, but also by playing $\rho_t = S$ and $a = \sigma(\emptyset)$, effectively starting the sequence $(a_0, a_1, ...)$ of Requirement 3. However, if player t hides the record and chooses $\rho_t = H$, then the strategy profile $\tilde{\sigma}$ implies that record keeping will start in period $t+1$ and the actions $(a_0, a_1, ...)$ will unfold from period $t+1$ instead. If Requirement 3 fails, there exists an action $\tilde{a}$ such that $V(\tilde{a}, a_0, a_1., ...) > V(a_0, a_1, ...) = \bar{V}$, which yields a higher payoff than following $\tilde{\sigma}$; this would imply that $\tilde{\sigma}$ is not an equilibrium strategy profile. □

We are now ready to prove Proposition 2.

Proof of Proposition 2. In the game in which record-keeping starts in period 0, let σ be a strategy profile whose equilibrium path is an organizational equilibrium. By Lemma 3 we can find a state-independent strategy profile $\tilde{\sigma}$ for the game in which history can be hidden that attains the same equilibrium path from $\hat{t}$ on, whatever the realization of $\hat{t}$; this equilibrium is also symmetric. To complete the proof, we need to show that there is no other state-independent equilibrium which is symmetric from period $\hat{t}$ on and attains a higher payoff from that point onwards. By contradiction, suppose that such an equilibrium existed, let it be $\bar{\tilde{\sigma}}$. From Lemma 1, we can assume without loss of generality that $\bar{\tilde{\sigma}}$ involves full revelation. Lemma 2 implies that $\bar{\tilde{\sigma}}|_{h^{\hat{t}}}$ is a symmetric state-independent equilibrium of the game in which history can be hidden, which would then achieve a higher payoff than σ; however, this would imply that σ does not satisfy Requirement 2 and therefore that its equilibrium path is not an organizational equilibrium, establishing a contradiction. □

C Further Discussion of Assumptions 1, 3, 4, and 5.

Assumption 1 is central to our definition. By ensuring that the preference ordering over sequences of actions is independent of the state, it provides a way of achieving a meaningful comparison across different periods of time (or different histories) for which the state variable is different. Section III.E provides an example where this assumption fails and illustrates a way we construct an approximating economy that satisfies it. Without uncertainty, utility functions are only identified up to monotone transformations. In this case, it can be shown that Assumptions 1 and 4 are equivalent. However, in the game of Section III.B, uncertainty is present, and we need the separability property to apply to *lotteries* about future outcomes. In this case, utility functions are identified up to affine transformations, and Assumption 4 is stronger than Assumption 1. Nonetheless, all of the separable preferences that we use in practice satisfy it. A (contrived) example of preferences that satisfies Assumption 1 but not Assumption 4 is one in which we amend the preferences of Section II to be

$$E_t\left[u(c_t)+\delta\sum_{\tau=1}^{\infty}\beta^{\tau}u\left(c_{t+\tau}\right)\right]^{\xi},$$

with $\xi < 1$: in addition to the standard risk aversion period by period (embedded in u), these preferences exhibit risk aversion over the entire infinite sequence. When Assumption 1 holds but Assumption 4 fails, an organizational equilibrium still exists, but the interpretation based on the alternative game of Section III.B does not necessarily apply. An avenue to generalize the results to this case would be to study the limiting behavior of the game of Section III.B to the probability of record-keeping being available in each period converging to 1.

The other Assumptions that we introduce are sufficient conditions that allow us to derive our results in a clean way, but there is often an alternative way to derive similar results in economies that do not satisfy them, in particular by relying on Proposition 3.

Specifically, we use Assumption 3 to prove that the game that includes only the action payoff $V(a_t, a_{t+1}, ...)$ has a reconsideration-proof equilibrium that satisfies the no-restarting condition. Weak separability between the initial action and the following sequence of actions allows us to find a worst continuation sequence that is a sufficient deterrent for all possible deviations. When Assumption 3 fails, the worst continuation may depend on the action taken, so for example the threat of restarting might work for the action to be taken in period $t+1$, but not in period $t+2$. Nonetheless, checking whether this is the case in an application is not a difficult exercise. As an example, consider the following modification of preferences and technology of Section II, that induce a violation of Assumption 3. At time t, preferences in terms of consumption sequences are given by

$$\frac{1}{1-\sigma}\left[c_t^{(1-\iota)(1-\sigma)}+\beta\delta\sum_{v=0}^{\infty}\beta^v\left(\frac{c_{t+v+1}}{c_{t+v}^{\iota}}\right)^{(1-\sigma)}\right],$$

with $\sigma \neq 1$ and $\iota \in (0,1]$,[8]

$$k_{t+1} = Ak_t - c_t.$$

Compared to the standard case, continuation preferences embed habit formation.[9] When we express this problem as preferences over a sequence of saving rates, so as to isolate the role of capital, we obtain

$$\frac{(Ak_t)^{(1-\iota)(1-\sigma)}}{1-\sigma}\left[(1-s_t)^{(1-\iota)(1-\sigma)}+\beta\delta\sum_{v=0}^{\infty}\beta^v\left(A^{v(1-\iota)+1}s_{t+v}(1-s_{t+v+1})(1-s_{t+v})^{-\iota}\prod_{n=0}^{v-1}s_{t+n}^{1-\iota}\right)^{1-\sigma}\right].$$

For these preferences, the marginal rate of substitution between s_{t+1} and s_{t+2} depends on s_t, so that separability of s_t from the remaining sequence fails. Nonetheless, we can establish whether an organizational equilibrium exists by computing it from a recursive structure. Even when Assumption 3 fails, the proof of Theorem 1 implies that a reconsideration-proof equilibrium of the game where preferences are given by the action component only exists. On the path of play implied by such an equilibrium, the value from the sequence of actions (excluding the separable state) is constant:

$$\frac{1}{1-\sigma}\left[(1-s_t)^{(1-\iota)(1-\sigma)}+\beta\delta\sum_{v=0}^{\infty}\beta^v\left(A^{v(1-\iota)+1}s_{t+v}(1-s_{t+v+1})(1-s_{t+v})^{-\iota}\prod_{n=0}^{v-1}s_{t+n}^{1-\iota}\right)^{1-\sigma}\right]=\bar{V}.$$

Using the fact that the players at time t and $t+1$ attain the same value $\bar{V}$, we derive a recursive expression similar to the one we derive in the applications of the main text:

$$\begin{aligned}\bar{V}(1-\beta(As_t)^{(1-\iota)(1-\sigma)}) &= (1-s_t)^{(1-\iota)(1-\sigma)}\\ &\quad+\beta\delta(As_t(1-s_{t+1})(1-s_t)^{-\iota})^{1-\sigma}-\beta(As_t(1-s_{t+1}))^{(1-\iota)(1-\sigma)} \qquad (3)\end{aligned}$$

For any given $\bar{V}$, equation (3) is a difference equation in the saving rates that can be solved numerically. For any given parameter combination, we can check whether this difference equation implies monotonic convergence to a steady state. As long as δ and ι are such that there is an incentive to undersave in the first period, that was the case in the numerical examples we tried. When this is the case, we can proceed as in the main text:

- Find the steady state that maximizes the value $\bar{V}$;
- From equation (3), derive the function that maps s_t into s_{t+1};

[8]When $\sigma = 1$ we obtain the logarithmic case, that preserves Assumption 3 even with the habit-formation specification here. $\iota = 0$ is the standard case in which Assumption 3 also applies.

[9]Introducing habit formation over the initial time-t consumption would break separability between the state and the actions.

- For any potential initial starting point s_0, we can compute the payoff that a player at time t receives if she expects restarting from s_0 to happen in period $t+1$ and plays the best response to it, and compare it to $\bar{V}$. If a value s_0 can be found such that the threat of reversion to s_0 in the future is enough to (weakly) deter any action, we have found an organizational equilibrium. Such a value for s_0 is guaranteed to exist under Assumption 3, and not here. Nonetheless, in the numerical examples we tried, there is an interval of values of s_0 where the condition is satisfied, just as in our applications in the main text, so an organizational equilibrium exists; as in the main text, we pick the highest saving in the interval where no-restarting applies based on Pareto optimality (though another choice would also be valid and converge to the same constant saving rate in the long run).

Finally, we used Assumption 5 to guarantee the existence of an organizational equilibrium which is recursive in the continuation value. As always in infinite-horizon models, a recursive structure is of great help for computations. Assumption 5 implies that the preference disagreement between the players moving at t and $t+1$ only concerns the action taken at $t+1$: conditional on the action taken at $t+1$, they agree on their preference ordering over sequences of actions from $t+2$ on. This allows us to use the continuation value $\widehat{V}(a_{t+2}, a_{t+3}, ...)$ as a state in computing the equilibrium path recursively. Even when Assumption 5 fails, there may be other ways of obtaining a recursive representation. As an example, we consider here a variant of the consumption-saving problem of Section II. We now assume that the planner is seeking to maximize the utility of a two-person household where both members have standard time-consistent preferences and share consumption, but they differ in their discount factor, generating time-inconsistency for the planner as in Jackson and Yariv (2014, 2015). Preferences at time t are thus given by

$$\sum_{v=0}^{\infty}(\beta_h^v + \lambda\beta_\ell^v)\log\left(c_{t+v}\right),$$

with $0 < \beta_\ell < \beta_h < 1$, and $\lambda > 0$ being a measure of the relative Pareto weight of the impatient member. Section II is a limiting case of these preferences as $\beta_\ell = 0$, $\beta_h = \beta$, and $\delta = 1/(1+\lambda)$. When $\beta_\ell > 0$, Assumption 5 fails, as we can see considering the relative discount factor between periods $t+2$ and $t+3$. From the perspective of period t, the relative discount factor is $(\beta_h^3 + \beta_\ell^3)/(\beta_h^2 + \beta_\ell^2)$, while from the perspective of period $t+1$ it is $(\beta_h^2 + \beta_\ell^2)/(\beta_h + \beta_\ell)$. As a consequence, the players at t and $t+1$ differ not only in the relative valuation of saving in period $t+1$, but also on saving in any future period. We can nonetheless retrieve a recursive structure for this game as well. Specifically, let $V_{\ell,t}(s_t, s_{t+1}, ...)$ and $V_{h,t}(s_t, s_{t+1}, ...)$ be the values accruing to the impatient and the patient member of the household respectively, when the planner chooses a sequence of savng rates $(s_t, s_{t+1}, ...)$, excluding the additive utility from initial capital $\alpha/(1-\alpha\beta_i)k_t$ for $i = h, \ell$. Since each member has standard time-consistent preferences, we can express these values recursively:

$$V_{i,t} = \log(1-s_t) + \frac{\alpha\beta_i}{1-\alpha\beta_i}\log s_t + \beta_i V_{i,t+1}, \quad i = h, \ell. \tag{4}$$

Since this economy satisfies the conditions of Proposition 1, an organizational equilibrium exists. The utility attained by the planner in such an equilibrium is a constant value $\bar{V} = V_{h,t} + \lambda V_{\ell,t}$. From the Pareto weighted sum of the two equations we then obtain

$$(1+\lambda)\log(1-s_t)+\left(\frac{\alpha\beta_h}{1-\alpha\beta_h}+\frac{\alpha\beta_\ell}{1-\alpha\beta_\ell}\right)\log s_t-\lambda(\beta_h-\beta_\ell)V_{\ell_{t+1}}=\bar{V}(1-\beta_h). \tag{5}$$

For any given value $\bar{V}$, equation (5) admits a unique solution for s_t as a function of $V_{\ell,t+1}$;[10] we can substitute this solution into (4) for $i=\ell$ and obtain a difference equation in $V_{\ell,t}$. We have thus obtained a recursive representation in terms of the value $V_{\ell,t}$. This difference equation can alternatively be expressed in terms of s_t, since (5) implies a one-to-one correspondence. In our numerical evaluations, this difference equation behaves in the same way as it does in the baseline case of Section II, so that the same procedure described above for the habit-formation economy can be used again here to compute the organizational equilibrium. Figure C.1 plots such an example for the transition function between s_t and s_{t+1}. When $\beta_L=0$, this economy becomes the standard quasi-hyperbolic discounting example in our baseline analysis. When $\beta_L>0$, the dynamics needs to be computed based on equations (4) and (5). In the end, the transitional dynamics are similar qualitatively.

Figure C.1: Evolution of the Saving Rate in Jackson and Yariv (2014)

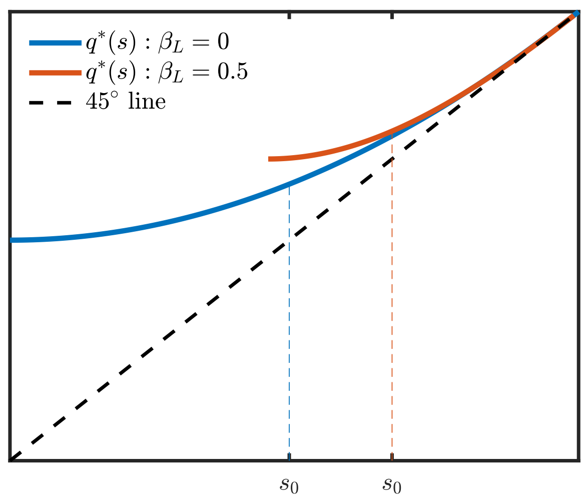

D Proofs of Section III.C.

D.1 Proof of Proposition 3

Proof. Suppose first that $\{\bar{a}_t\}_{t=0}^{\infty}$ is a sequence satisfying the three properties in the proposition. We construct a subgame-perfect equilibrium strategy profile as follows.[11] We start with $\sigma_0(\emptyset)=\bar{a}_0$. Let $\tilde{h}^t$, $t\geq 1$ be an

[10] More precisely, the equation admits at most one solution, and may have none. However, since we know that an organizational equilibrium exists, a solution has to exist for the appropriate range of values of $\bar{V}$ and V_ℓ.

[11] We defined an organizational equilibrium within the context of the game of Section III.A, so the proposition is proven in the context of this game, although of course the results apply to the game of Section III.B when Assumption 4

arbitrary history whose predecessors are $(\emptyset, \tilde{h}^0, \tilde{h}^1, ...\tilde{h}^{t-1})$. If $a_s = \sigma_s(\tilde{h}^{s-1})$, $s = 0, ..., t-1$, set $\sigma_t(\tilde{h}^t) = \bar{a}_t$. Otherwise, let $\tilde{t} := \max\{s : a_s \neq \sigma_s(\tilde{h}^{s-1})\}$ and set $\sigma_t(\tilde{h}^t) = \bar{a}_{t-1-\tilde{t}}$. In words, this strategy punishes any deviation by restarting the continuation equilibrium from the same equilibrium path that is supposed to prevail in period 0. Properties 1 and 3 ensure that such a punishment is sufficient to deter deviations, both in the initial period and in any subsequent period and history. This equilibrium is state independent (Requirement 1) and symmetric, since the equilibrium path of play attains an action value $\bar{V}$ independent of the past history. No equilibrium can attain a higher constant value. Suppose such an equilibrium existed, and let $\{a_t^B\}_{t=0}^{\infty}$ be its equilibrium path, which attains a constant $V^B > \bar{V}$. Then we would have $V(a_t^B, a_{t+1}^B, ...) = V^B > \bar{V}$, $\forall t \geq 0$, which would contradict property 2 of our initial sequence. Therefore, the newly constructed subgame-perfect equilibrium satisfies Requirement 2. Finally, Requirement 3 is a direct analog of the third property that we imposed on the sequence.

Suppose now that a sequence satisfying the 3 properties of the proposition exists and its value is $\bar{V}$. Requirement 2 implies that all organizational equilibria feature a path of constant value $\bar{V}$ as well, which implies that they satisfy the first two properties; the third property follows directly from Requirement 3. □

D.2 Proof of Proposition 4

To prove this we rely on a useful lemma, which introduces a convenient way of representing equilibria through their values, similarly to Abreu, Pierce, and Stacchetti's (1986; 1990) method.[12]

Lemma 4. *Let $V^* \in \mathbb{R}$ and $\hat{\mathcal{V}} \subset \mathbb{R}$ be a value and a set of continuation values that satisfy the following properties:*

1. $$\forall a \in A \quad \exists \hat{v} \in \hat{\mathcal{V}} : \tilde{V}(a, \hat{v}) \leq V^*;$$

2. $$\forall v \in \hat{\mathcal{V}} \quad \exists (a, \hat{v}) \in A \times \hat{\mathcal{V}} : \tilde{V}(a, \hat{v}) = V^* \wedge W(a, \hat{v}) = v.$$

3. *There exists no value $V^{**} > V^*$ and set $\hat{\hat{\mathcal{V}}}$ that satisfies properties 1 and 2; furthermore, there is no set $\hat{\mathcal{V}}_a \supset \hat{\mathcal{V}}$ that satisfies properties 1 and 2 together with V^*.*

Then:

- *Construct an arbitrary sequence of actions $\{a_t^*\}_{t=0}^{\infty}$ recursively as follows. In period 0, pick $\hat{v}_0^* \in \hat{\mathcal{V}}$ and $(a_0^*, \hat{v}_1^*) \in A \times \hat{\mathcal{V}}$ such that $\tilde{V}(a_0^*, \hat{v}_1^*) = V^*$ and $W(a_0^*, \hat{v}_1^*) = \hat{v}_0^*$. In each subsequent period, pick*

is satisfied.

[12]Note, however, that we cannot adopt their method to recursively compute the desired sets. *Given* V^*, $\hat{\mathcal{V}}$ can be computed recursively as in Abreu, Pierce, and Stacchetti. However, without further assumptions the set of values of V^* for which $\hat{\mathcal{V}}$ is defined need not be convex, which makes finding its maximum difficult.

$(a_t^*, \hat{v}_{t+1}^*) \in A \times \hat{\mathcal{V}}$ such that $\tilde{V}(a_t^*, \hat{v}_{t+1}^*) = V^*$ and $W(a_t^*, \hat{v}_{t+1}^*) = \hat{v}_t^*$. Constructing such a sequence is possible by the definition of V^* and $\hat{\mathcal{V}}$. The sequence so constructed is the outcome of a reconsideration-proof equilibrium;

- *If $\{a_t^*\}_{t=0}^{\infty}$ is the equilibrium path of a reconsideration-proof equilibrium, $\tilde{V}(a_0^*, a_1^*, ...) = V^*$ and $\hat{V}(a_t^*, a_{t+1}^*, ...) \in \hat{\mathcal{V}}$ for any $t > 0$.*

Proof.

First, we prove that the recursively-constructed sequence $\{a_t^*\}_{t=0}^{\infty}$ satisfies

$$\tilde{V}(a_t^*, \hat{V}(a_{t+1}^*, a_{t+2}^*, ...)) = V^* \quad \forall t \geq 0 \tag{6}$$

and

$$\widehat{V}(a_t^*, a_{t+1}^*, a_{t+2}^*, ...) \in \hat{\mathcal{V}} \quad \forall t \geq 0. \tag{7}$$

Note that, if $\widehat{v}_T^* = \widehat{V}(a_T^*, a_{T+1}^*, a_{T+2}^*, ...)$ for some period T, iterating backwards we find that $\widehat{v}_t^* = \widehat{V}(a_t^*, a_{t+1}^*, a_{t+2}^*, ...)$ for all $t < T$, so that equations (6) and (7) hold.

Define

$$\{\underline{a}_t\}_{t=0}^{\infty} \in \arg\min_{\{a_t\}_{t=0}^{\infty}} \widehat{V}(a_0, a_1, ...)$$

and similarly let $\{\bar{a}_t\}_{t=0}^{\infty}$ be a sequence that attains the maximum. Both exist by the compactness of A and the continuity of $\widehat{V}$ (in the product topology).

Next, truncate the sequence $\{a_t^*\}_{t=0}^{\infty}$ at time $S > T$ and replace the continuation with $\{\underline{a}_t\}_{t=0}^{\infty}$ or $\{\bar{a}_t\}_{t=0}^{\infty}$. By Assumption 5 and the monotonicity of W, we have

$$\widehat{V}(a_T^*, a_{T+1}^*, ..., a_S^*, \underline{a}_0, \underline{a}_1, ...) \leq \widehat{V}(a_T^*, a_{T+1}^*, ..., a_S^*, a_{S+1}^*, a_{S+2}^*, ...) \leq \widehat{V}(a_T^*, a_{T+1}^*, ..., a_S^*, \bar{a}_0, \bar{a}_1, ...) \tag{8}$$

and

$$\begin{aligned}
&\widehat{V}(a_T^*, a_{T+1}^*, ..., a_S^*, \underline{a}_0, \underline{a}_1, ...) = W(a_T^*, W(a_{T+1}^*, ...W(a_S^*, W(\underline{a}_0, W(\underline{a}_1, ...)...))...)) \leq \\
&W(a_T^*, W(a_{T+1}^*, ...W(a_S^*, \hat{v}_S^*)...)) = \hat{v}_T^* \leq \\
&W(a_T^*, W(a_{T+1}^*, ...W(a_S^*, W(\bar{a}_0, W(\bar{a}_1, ...)...))...)) = \widehat{V}(a_T^*, a_{T+1}^*, ..., a_S^*, \bar{a}_0, \bar{a}_1, ...).
\end{aligned} \tag{9}$$

Taking limits as $S \to \infty$ in equations (8) and (9) and exploiting the continuity of $\widehat{V}$ according to the product topology, the left-most and right-most expressions in the inequalities converge to the same value, which then implies that indeed $\widehat{v}_T^* = \widehat{V}(a_{T+1}^*, a_{T+2}^*, a_{T+3}^*, ...)$ and (6) and (7) hold.

To complete the proof of the first point, we need to show that there exists no symmetric subgame-perfect equilibrium whose payoff is strictly greater than V^*. By contradiction, suppose that there is such an equilibrium with value $V^{**} > V^*$. Let σ^{**} be the strategy profile representing one such equilibrium. Define

$$\hat{\mathcal{V}}_b := \{v : v = \widehat{V}(a^{**}_{t+1}|_{h^t}, a^{**}_{t+2}|_{h^t}, a^{**}_{t+3}|_{h^t}, ...),\, h^t \in A^t\},$$

where $\{a^{**}_s|_{h^t}\}^{\infty}_{s=t+1}$ is the equilibrium path implied by the strategy profile σ^{**} following a history h^t. The pair $(V^{**}, \hat{\mathcal{V}}_b)$ satisfies property 1 in the lemma, since otherwise σ^{**}_0 would not be optimal at time 0. It also satisfies property 2 since σ^{**} is symmetric and by the definition of $\hat{\mathcal{V}}_b$. But then this implies that property 3 in the lemma does not hold for V^*, establishing a contradiction.

In the previous point we proved that, given V^* and $\hat{\mathcal{V}}$, we can construct a reconsideration-proof equilibrium of value V^*. Since all reconsideration-proof equilibria must have the same value, it must be the case that $\tilde{V}(a^*_0, a^*_1, ...) = V^*$. Furthermore, repeating the steps of the previous point, we can prove that the value V^* and the set

$$\hat{\mathcal{V}}_a := \{v : v = \widehat{V}(a^{*}_{t+1}|_{h^t}, a^{*}_{t+2}|_{h^t}, a^{*}_{t+3}|_{h^t}, ...),\, h^t \in A^t\}$$

satisfy properties 1 and 2. By the definition of $\hat{\mathcal{V}}$, it follows that $\hat{\mathcal{V}}_a \subseteq \hat{\mathcal{V}}$. □

While not essential for the proof of Proposition 4, the following lemma is useful for computations:

Lemma 5. *The set $\hat{\mathcal{V}}$ defined in Lemma 4 is convex.*[13]

Proof. We first define the set $\hat{\mathcal{V}}_c$ by relaxing property 2 in Lemma 4 to be the following:

$$\forall v \in \hat{\mathcal{V}}_c \quad \exists (a, \hat{v}) \in A \times \hat{\mathcal{V}} : \tilde{V}(a, \hat{v}) \geq V^* \wedge W(a, \hat{v}) = v. \tag{10}$$

We will later prove that $\hat{\mathcal{V}}_c = \hat{\mathcal{V}}$.

Simple case. First, if $\hat{\mathcal{V}}_c$ is a singleton, then it is necessarily convex and $\hat{\mathcal{V}}_c = \hat{\mathcal{V}}$: by property 3 of Lemma 4, V^* should be raised until $\tilde{V}(a, \hat{v}) = V^*$ at the single element $\hat{v} \in \hat{\mathcal{V}}_c$, with no effect on property 2 and relaxing the constraint in property 1.

From now on, we study the case in which $\hat{\mathcal{V}}_c$ contains at least two values.

Step 1. To prove that $\hat{\mathcal{V}}_c$ is convex, we prove that its convex hull, $\text{Co}(\hat{\mathcal{V}}_c)$, satisfies properties 1 and 2 as well (and of course $\text{Co}(\hat{\mathcal{V}}_c) \supset \hat{\mathcal{V}}_c$ unless $\hat{\mathcal{V}}_c$ is convex as well). Property 1 is immediate from the monotonicity of $\tilde{V}$. Let $v_1, v_2 \in \hat{\mathcal{V}}_c$, and let $(a_1, \hat{v}_1)$, $(a_2, \hat{v}_2)$ elements of $A \times \hat{\mathcal{V}}_c$ be two pairs of actions and continuation values that satisfy property 2 of Lemma 4. Consider their convex combination $(\alpha v_1 + (1-\alpha) v_2, \alpha \hat{v}_1 + (1-\alpha)\hat{v}_2)$, $\alpha \in [0, 1]$. Since $\tilde{V}$ is continuous and quasiconcave and W is continuous, $\tilde{V}(\alpha v_1 + (1-\alpha)v_2, \alpha\hat{v}_1 + (1-\alpha)\hat{v}_2) \geq V^*$, and $W(\alpha v_1 + (1-\alpha)v_2, \alpha\hat{v}_1 + (1-\alpha)\hat{v}_2)$ takes all values in $[v_1, v_2]$ as α varies between 0 and 1. Hence, all intermediate values satisfy property 2 as well, which completes the proof that $\text{Co}(\hat{\mathcal{V}}_c)$ satisfies property 2.

[13] Lemma 4 defines a unique set, since the union of all sets satisfying properties 1 and 2 satisfies properties 1 and 2 as well.

Step 2. To prove that $\hat{\mathcal{V}}_c = \hat{\mathcal{V}}$, proceed as follows. Define $\underline{v}_c := \min\{\hat{\mathcal{V}}_c\}$ and $\bar{v}_c := \max\{\hat{\mathcal{V}}_c\}$.[14] By definition, we can find $(\underline{a}, \hat{\underline{v}})$ and $(\bar{a}, \bar{\hat{v}})$ such that

$$\tilde{V}(\underline{a}, \hat{\underline{v}}) \geq V^* \wedge W(\underline{a}, \hat{\underline{v}}) = \underline{v}$$

and

$$\tilde{V}(\bar{a}, \bar{\hat{v}}) \geq V^* \wedge W(\bar{a}, \bar{\hat{v}}) = \bar{v}.$$

Since A is convex, we can construct within it a line from $\underline{a}$ to $\bar{a}$ by defining $a(\alpha) := \alpha\underline{a} + (1-\alpha)\bar{a}$, $\alpha \in [0,1]$. By the quasiconcavity of $\tilde{V}$, we know

$$\tilde{V}(a(\alpha), \alpha\hat{\underline{v}} + (1-\alpha)\bar{\hat{v}}) \geq V^*.$$

By property 1 of Lemma 4, for each action $a(\alpha)$ and the monotonicity and continuity of $\tilde{V}$ we have

$$\tilde{V}(a(\alpha), \underline{v}) \leq V^*.$$

Since $\hat{\mathcal{V}}_c$ is convex, we can find a (unique) value $\hat{v}(\alpha)$ such that

$$\tilde{V}(a(\alpha), \hat{v}(\alpha)) = V^*.$$

Monotonicity and continuity of $\tilde{V}$ imply that $\hat{v}(\alpha)$ is a continuous function. It then follows that $\hat{V}(a(\alpha), \hat{v}(\alpha))$ is a continuous function of α. As $\alpha \in [0,1]$, this function must take all values between $\underline{v}$ and $\bar{v}$, proving that the property 2 of Lemma 4 is satisfied by $\hat{\mathcal{V}}_c$ and thus $\hat{\mathcal{V}}_c = \hat{\mathcal{V}}$. □

We are now ready to prove Proposition 4.

Proof. The second property of the value V^* and the set $\hat{V}$ in Lemma 4 implies that we can construct a function $g : \hat{\mathcal{V}} \to \mathbb{R} \times \hat{\mathcal{V}}$ with the property that $\tilde{V}(g(v)) = V^*$ and $W(g(v)) = v$. Starting from any value $v_0 \in \hat{\mathcal{V}}$, we can construct recursively a path $(a_t, v_{t+1}) = g(v_t)$. By Lemma 4, this is the equilibrium path of a reconsideration-proof equilibrium. It will thus be an organizational equilibrium provided that

$$V(a_t, v_{t+1}) \geq \max_a \tilde{V}(a, v_0) \quad \forall t.$$

By the definition of $\mathcal{V}$, this property is satisfied by its least element, $\underline{v}$;[15] hence, it will be satisfied provided that the initial value v_0 is sufficiently low. □

[14] It is straightforward to prove that $\hat{\mathcal{V}}_c$ is closed, by the continuity of the functions defining it.

[15] By the monotonicity of $\tilde{V}$ in its second argument and the property 1 of $\mathcal{V}$, $\tilde{V}(a, \underline{v}) \leq V^*$ for all $a \in A$.

D.3 Proof of Proposition 5

Proof. Define a correspondence $\zeta : \mathbb{R} \times \mathbb{R} \rightrightarrows \mathbb{R}$ as follows:

$$v \in \zeta(v', v^*) \Longleftrightarrow \exists a \in A : \begin{cases} \tilde{V}(a, v') = v^* \\ W(a, v') = v. \end{cases} \tag{11}$$

In words, given (v^*, v'), v belongs to the correspondence if there is an action a which, together with a continuation value v', yields utility v^* when evaluated according to the decision maker's preferences $(\tilde{V})$ and utility v when evaluated with its continuation utility function W.

We prove that there exists a value v^* for which ζ is nonempty and admits a fixed point in continuation utilities $(v = v')$. We do so by proving that a Markov equilibium (a^M, v^M) exists, such that

$$v^* = \tilde{V}(a^M, v^M) = \max_a \tilde{V}(a, v^M) \tag{12}$$

and

$$v^M = W(a^M, v^M). \tag{13}$$

To prove the existence of a Markov equilibrium, we construct a correspondence $\hat{a}(.)$ from A into itself by setting

$$\hat{a}(a) = \max_{a_0 \in A} \widehat{V}(a_0, a, a, a, ...).$$

By the usual compactness and continuity properties, this correspondence is nonempty, compact-valued, and upper hemicontinuous. Quasiconcavity of $\widehat{V}$ ensures that it is also convex-valued. Hence, the correspondence has a fixed point by Kakutani's theorem; let a^M be one such fixed point. Given Assumption 5, letting $v^M := \widehat{V}(a^M, a^M, a^M, ...)$, equations (12) and (13) are satisfied.

We thus know $v^M \in \zeta(v^M, \tilde{V}(a^M, v^M))$. Once again, our assumptions about compactness and continuity imply that the correspondence ζ is upper hemicontinuous. Let V^* be the maximal value for which ζ admits a fixed point in continuation utilities. In the proofs below, it is useful to establish that

$$v \in \zeta(v', V^*) \Longrightarrow v \leq v'. \tag{14}$$

Suppose (14) is not satisfied. Let (a, v') be such that $V(a, v') = V^*$ and $W(a, v') > v'$. Holding the action a fixed, continuity and monotonicity imply that higher values of v' lead to higher values of $V(a, v')$ and $W(a, v')$. As long as $W(a, v') > v'$, we know that $v' < \max_{\{a_t\}_{t=0}^{\infty}} \widehat{V}(a_0, a_1, ...)$ and can thus be raised further. Eventually, we will attain a value $v^h > v'$ for which $W(a, v^h) = v^h$ (this has to happen, since $W(a, v')$ is bounded by the maximum above). Let $V^h := V(a, v^h) > V^*$. We just established that a fixed point of $\zeta(., V^h)$ exists, which contradicts the assumption that V^* is the highest value for which a fixed point can be found.

In our next step, we prove that there are no symmetric equilibria with value $V^{**} > V^*$. By the definition of V^*, given any combination of an action and a continuation utility (a, v'), if $\tilde{V}(a, v') = V^{**}$ then $W(a, v') < v'$. This implies that any equilibrium path with value V^{**} would feature a strictly increasing sequence of continuation values; convergence is ruled out, because continuity and compactness would imply that the limiting point would be a fixed point of ζ, which is inconsistent with $V^{**} > V^*$. Since the set of possible continuation values is bounded by

$$\max_{\{a_t\}_{t=0}^{\infty}} \widehat{V}(a_0, a_1, \ldots),$$

no such equilibrium path can exist.

We now prove that there exist symmetric equilibria with value V^*, which then implies that any such equilibrium is reconsideration proof. Let v^{SS} be the maximal fixed point of $\zeta(., V^*)$. For any continuation value $v > v^{SS}$, a repetition of the arguments described above for V^{**} imply that no equilibrium path would be possible.[16] We prove instead that there exists a convex set $\mathcal{V} = [v_\ell, v^{SS}]$ which, together with V^*, satisfies the properties of Lemma 4, where

$$v_\ell := \min_{v' \leq v^{SS}} \min \zeta(v', V^*). \tag{15}$$

To do so, prove first that, for any action $a \in A$, $\tilde{V}(a, \min_{\{a_t\}_{t=0}^{\infty}} \widehat{V}(a_0, a_1, ...)) \leq V^*$. By contradiction, suppose that an action a_L such that $\tilde{V}(a_L, \min_{\{a_t\}_{t=0}^{\infty}} \widehat{V}(a_0, a_1, ...)) > V^*$ existed. We could then repeat the same steps used to prove (14) and construct a steady state with value higher than V^*.

Since $\tilde{V}(a, \min_{\{a_t\}_{t=0}^{\infty}} \widehat{V}(a_0, a_1, ...)) \leq V^* \quad \forall a \in A$, we can define

$$v'_{\min} := \min_{(a,v')} v' := \tilde{V}(a, v') = V^*.$$

Since there exists an action a^{SS} such that $V(a^{SS}, v^{SS}) = V^*$, $v'_{\min} \leq v^{SS}$. Also, by equations (14) and (15), $v_\ell \leq v'_{\min}$. Hence, $\tilde{V}(a, v_\ell) \leq V^* \quad \forall a \in A$: Property 1 of Lemma 4 is satisfied by the value V^* and the continuation set $[v_\ell, v^{SS}]$. To prove Property 2, let a_ℓ and v'_ℓ be such that $W(a_\ell, v'_\ell) = v_\ell$ and $\tilde{V}(a_\ell, v'_\ell) = V^*$, and $\lambda \in [0, 1]$.[17] As we just established, $\tilde{V}(\lambda a_\ell + (1-\lambda)a^{SS}, v_\ell) \leq V^*$. By quasiconcavity, $\tilde{V}(\lambda a_\ell + (1-\lambda)a^{SS}, \lambda v'_\ell + (1-\lambda)v^{SS}) \geq V^*$. Strict monotonicity implies that there exists a unique value v_λ such that $\tilde{V}(\lambda a^{SS} + (1-\lambda)a_\ell, v_\lambda) = V^*$, which must vary continuously with λ by the continuity of $\tilde{V}$. It follows that $W(\lambda a^{SS} + (1-\lambda)a_\ell, v_\lambda)$ is a continuous function of λ and it takes all values between v_ℓ and v^{SS}, proving that Property 2 of Lemma 4 holds. Finally, from equations (14) and (15), we know that any value $v \notin [v_\ell, v^{SS}]$ could only be attained by some action a with a continuation value $v' > v^{SS}$, which would lead to nonexistence in subsequent periods. Hence, $[v_\ell, v^{SS}]$ is the largest set that satisfies Properties 1 and 2 of Lemma 4 together with the value V^*, completing the proof that a reconsideration-proof equilibrium has value V^*, and thus that in turn the organizational equilibrium with the state variable is also associated with an action value V^*. Our construction also proved that V^* is the maximal action payoff that can be attained by a constant action.

Finally, suppose that $\widehat{V}$ is strictly quasiconcave. Let a^{SS} be the unique action that attains $\max_a V(a, a, a, ...)$. If this steady state is not a Markov equilibrium, then $a^{SS} < \max_a \tilde{V}(a, v^{SS})$. In this case, a sequence that

[16] If along the equilibrium path, for some $T \geq 0$, $v_T > v^{SS}$, then $v_t > v^{SS}$ for all $t > T$. Since $\{v_t\}$ is bounded and monotonically increasing, the limiting point will be a fixed point of ζ, which is incompatible with v^{SS} being the largest fixed point.

[17] We have $v_\ell \leq v'_\ell \leq v^{SS}$ by (14) and (15).

starts at a^{SS} and stays constant violates the no-delay condition.

Finally, we prove part 2 of the proposition. The Ramsey outcome is the allocation that attains the highest payoff, and so by definition an organizational equilibrium cannot do better. If there is no constant allocation that attains the Ramsey outcome, then it means that the best constant allocation attains a payoff strictly smaller than Ramsey; Proposition 5 proves that the payoff of an organizational equilibrium coincides with that of the best constant allocation, and is thus strictly worse than Ramsey as well. When a constant allocation a^{SS} attains the Ramsey outcome, it must be the case that

$$a_{SS} \in \arg\max_a V(a, a^{SS}, a^{SS}, ...);$$

this implies that $(a^{SS}, a^{SS}, a^{SS}, ...)$ is also a Markov equilibrium, and that a^{SS} achieves the highest payoff among constant allocations, which (by Proposition 5) is also the payoff of an organizational equilibrium. In particular, $(a^{SS}, a^{SS}, ...)$ is an organizational equilibrium.

A state-independent Markov equilibrium cannot depend on the past nor on calendar time, and so it is a constant sequence $(a, a, ...)$. An organizational equilibrium attains the same payoff as the best constant allocation; hence, it can be no worse than the best Markov equilibrium, and is strictly better whenever the best constant allocations do not correspond to a Markov equilibrium.

□

D.4 Proof of Corollary 1

Proof. This proof follows closely that of Proposition 5. Let ζ, V^*, v_ℓ, and v^{SS} be defined as in that proof. The proof of Proposition 5 rules out symmetric equilibria with values higher than V^* by showing that there does not exist a sequence of actions that has a constant value higher than V^*. It also shows how to construct a sequence such that $\tilde{V}(a_0, \hat{V}(a_1, a_2, ...)) = V^*$ and $\hat{V}(a_0, a_1, ...) = v$ for any value in $v \in [v_\ell, v^{SS}]$; any such sequence satisfies properties 1 and 2 of Proposition 3. Let $\{\bar{a}_t\}_{t=0}^{\infty}$ be such that $\tilde{V}(\bar{a}_0, \hat{V}(\bar{a}_1, \bar{a}_2, ...)) = V^*$ and $\hat{V}(\bar{a}_0, \bar{a}_1, ...) = v_\ell$. The proof of Proposition 5 establishes that $\tilde{V}(a, v_\ell) \leq V^* \quad \forall a \in A$. Hence, $\{\bar{a}_t\}_{t=0}^{\infty}$ satisfies Property 3 of Proposition 3 as well. □

E Proofs for Section III.D

In this appendix, we establish that the slope of the transition function in the organizational equilibrium equals to 1 when approaching to the steady state and equals to 0 when starting at the saving rate in the Markov equilibrium. Furthermore, the slope is positive between the steady state and the Markov saving rate.

Given the transition function (13), the slope of it can be expressed as

$$\frac{\partial s_{t+1}}{\partial s_t} = -\exp\left\{\frac{-(1-\beta)V^* + \frac{\delta\alpha\beta}{1-\alpha\beta}\log s_t + \log(1-s_t)}{\beta(1-\delta)}\right\}\left(\frac{\frac{\delta\alpha\beta}{1-\alpha\beta}\frac{1}{s_t} - \frac{1}{1-s_t}}{\beta(1-\delta)}\right).$$

With a constant saving rate s, the lifetime action payoff is

$$(1-\beta)\overline{V} = \frac{\delta\alpha\beta}{1-\alpha\beta}\log s + \log(1-s) - \beta(1-\delta)\log(1-s).$$

The optimal constant saving rate s^* satisfies

$$\frac{\delta\alpha\beta}{1-\alpha\beta}\frac{1}{s^*} - \frac{1}{1-s^*} - \frac{\beta(1-\delta)}{1-s^*} = 0,$$

and the action payoff V^* satisfies

$$(1-\beta)V^* = \frac{\delta\alpha\beta}{1-\alpha\beta}\log s^* + \log(1-s^*) - \beta(1-\delta)\log(1-s^*).$$

Therefore, we have

$$\begin{aligned}
\left.\frac{\partial s_{t+1}}{\partial s_t}\right|_{s_t=s^*} &= -\exp\left\{\frac{-(1-\beta)V^* + \frac{\delta\alpha\beta}{1-\alpha\beta}\log s^* + \log(1-s^*)}{\beta(1-\delta)}\right\}\left(\frac{\frac{\delta\alpha\beta}{1-\alpha\beta}\frac{1}{s^*} - \frac{1}{1-s^*}}{\beta(1-\delta)}\right) \\
&= -(1-s^*)\left(\frac{\frac{\delta\alpha\beta}{1-\alpha\beta}\frac{1}{s^*} - \frac{1}{1-s^*}}{\beta(1-\delta)}\right) \\
&= 1
\end{aligned}$$

In the Markov equilibrium, the saving rate s^M maximizes the part involving only the current saving rate:

$$\frac{\delta\alpha\beta}{1-\alpha\beta}\log s + \log(1-s),$$

which implies that

$$\frac{\delta\alpha\beta}{1-\alpha\beta}\frac{1}{s^M} - \frac{1}{1-s^M} = 0.$$

As a result, $\left.\frac{\partial s_{t+1}}{\partial s_t}\right|_{s_t=s^M} = 0.$

Denote $\chi(s_t) \equiv \frac{\delta\alpha\beta}{1-\alpha\beta}\frac{1}{s_t} - \frac{1}{1-s_t}$. Notice that: (1) $\chi(s_t)$ is decreasing in s_t when $s_t \in (0,1)$; (2) $\chi(s_t) = 0$ when $s_t = s^M$. It follows that, $\frac{\partial s_{t+1}}{\partial s_t} > 0$ when $s_t > s^M$.

F Example of Approximating Strategy

As an example to illustrate the approximating strategy, we revisit the quasi-geometric discounting economy with partial depreciation and CRRA utility function and apply our approximation strategy. Compared with the environment in Section 2, we modify the period utility function and the law of motion to be

$$u(c) = \frac{c^{1-\sigma}}{1-\sigma}, \qquad k_{t+1} = f(k_t) - c_t + (1-d)k_t,$$

where $d \in (0, 1)$.

Let s_t denote the saving rate. Mapping to the general setup, we have

$$\begin{aligned}
P(k, s) &= u((1-s)f(k)), \\
Q(k, s) &= \beta(\delta - 1)u((1-s)f(k)), \\
F(k, s) &= sf(k) + (1-d)k.
\end{aligned}$$

In this economy, the action s_t and the states k_t are not separable. To proceed, we choose $m(s) = \frac{s^{1-\sigma}}{1-\sigma}$ to approximate the utility function and $g(s) = \log(1-s)$ to approximate the technology. In this approximating economy, the organizational equilibrium can be constructed.

The blue solid line in Figure displays the transition paths of the capital shock and the saving rate of the organizational equilibrium in this approximating economy. The red dashed line and the black broken line correspond to the Markov equilibrium and the Ramsey outcome in the approximating economy. Similar to our baseline analysis, the organizational equilibrium gradually transits from being close to the Markov equilibrium towards being close to the Ramsey outcome.[18]

To evaluate this approximation, we also compute the the Markov equilibrium and the Ramsey outcome in the original economy via global solutions, which are shown by the lines with circle markers. The outcomes in the approximating economy and the original economy are close to each other not only in the steady states but also along the entire transition paths.

G Organizational Equilibrium in Policy Problems

In Section III, there is one player for each period. Here, the policymaker is still represented by one player for each period, but we also include a continuum of identical households that face a dynamic problem.[19] In this appendix, we describe explicitly the strategic interaction between the government and the households at different points in time. The game unfolds as follows. In each period, the government in power takes an action $a \in A$ first. Then, the households move simultaneously. Each household takes an action $s \in S$. The aggregate state for next period evolves according to $k' = F(k, a, s)$. A full description would require us to specify what happens when households take different actions, so that, while they are identical ex ante, they may end up being different ex post. However, in most of the applications that are of interest, the household optimization

[18] We set $\beta = 0.8$, $\delta = 0.9$, $d = 0.5$, $\sigma = 2$, $\alpha = 0.36$.

[19] The notion of an equilibrium can be readily extended to environments with finite types of households or to economies with overlapping generations. Extending organizational equilibrium to economies with a continuum of types could be done by interacting the analysis here with distributional notions of equilibrium as in Jovanovic and Rosenthal (1988).

Figure F.1: Transition Paths in Approximating Economy

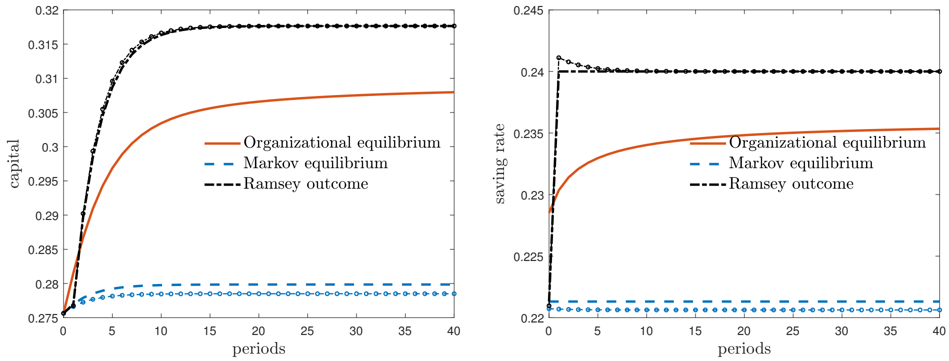

Note: The dotted lines in red and black represent the true solutions to the Markov equilibrium and the Ramsey outcome, respectively. The dashed lines in red and black represent the solutions in the approximating economy according to our strategy.

problem has a unique solution. Hence, there can be no equilibrium in which identical households take different actions. Moreover, a deviation by a single household has no effect on aggregates. We exploit these properties and specify the evolution of the economy and preferences only after histories in which (almost) all households have taken the same action. Starting from an arbitrary period t and state k_t, household preferences are given by a function

$$Z(k_t, \{a_v, s_v, s_v^-\}_{v=t}^{\infty}), \tag{16}$$

where s_v represents the action taken by the individual household, and s_v^- is the action taken by (almost) all other households. We assume that S is a convex compact subset of a locally convex topological linear space and that Z is jointly continuous in all of its arguments (in the product topology), strictly quasiconcave in the own action sequence $\{s_v\}_{v=t}^{\infty}$, and weakly separable between the state and the remaining arguments. We also assume that household preferences are time consistent. More precisely, we assume that, given an initial level of the state k_t and a sequence of other households' actions $\{a_v, s_v\}_{v=t}^{\infty}$,

$$\begin{aligned} &Z(k_t, \{a_v, s_v, s_v\}_{v=t}^{\infty}) = \max_{\{\tilde{s}_v\}_{v=t}^{\infty}} Z(k_t, \{a_v, \tilde{s}_v, s_v\}_{v=t}^{\infty}) \Longrightarrow Z(F(k_t, a_t, s_t), \{a_v, s_v, s_v\}_{v=t+1}^{\infty}) = \\ &\max_{\{\tilde{s}_v\}_{v=t+1}^{\infty}} Z(k_t, \{a_v, \tilde{s}_v, s_v\}_{v=t+1}^{\infty}). \end{aligned} \tag{17}$$

Equation (17) states that, if it is optimal from period t to follow the same sequence of actions that all other households are taking, then it is also optimal to follow that sequence in subsequent periods, as long as other households also continue to do the same. Notice that we exploit the fact that each household has no effect on the aggregates to leave the continuation preferences over several histories unspecified; this is convenient, because it prevents us from having to explicitly introduce individual state variables. To be concrete, consider the taxation game to which we apply this general definition; in that game, s_t is the individual saving rate. Equation (17) is written from the perspective of a household that starts with the same level of k_t as the aggregate, which allows us not to draw a distinction between the two. If that household finds it optimal to follow the same saving rate as all other households, then it will optimally choose to have the same level of

k_{t+1}, and equation (17) ensures that the continuation plan will remain optimal from period $t+1$ onwards. If instead the household chooses a different saving rate from others, then it would potentially enter period $t+1$ with a different level of the state from the aggregate; however, whenever this choice does not maximize (16), we know this would not be an optimal individual choice without need to specify the entire continuation path; moreover, the individual deviation does not affect aggregate incentives; hence, we do not need to keep track of it for the purpose of computing other households' best response either.

We define a competitive equilibrium from period t and a state k_t as a sequence $\{a_v, s_v\}_{v=t}^{\infty}$, such that

$$Z(k_t, \{a_v, s_v, s_v\}_{v=t}^{\infty}) = \max_{\{\tilde{s}_v\}_{v=t}^{\infty}} Z(k_t, \{a_v, \tilde{s}_v, s_v\}_{v=t}^{\infty}).$$

Proposition G.1. *Given any sequence of policy actions $\{a_v\}_{v=t}^{\infty}$, a competitive equilibrium exists.*

Proof. Fix k_t and $\{a_v\}_{v=t}^{\infty}$. Given our assumptions on S and Z, the best-response function

$$br(\{s_v\}_{v=0}^{\infty}) := \arg \max_{\{\tilde{s}_v\}_{v=t}^{\infty}} Z(k_t, \{a_v, \tilde{s}_v, s_v\}_{v=t}^{\infty})$$

is well defined and continuous. By Brouwer's theorem, it admits a fixed point, which is a competitive equilibrium. □

Equation (17) ensures that the continuation of a competitive equilibrium is a competitive equilibrium itself. Also, the separability assumption about Z implies that, if $\{a_v, s_v\}_{v=t}^{\infty}$ is a competitive equilibrium from a state k_t, then it is also a competitive equilibrium from any other state k'_t.

In what follows, we proceed by assuming that the competitive equilibrium is unique given a sequence of policy actions, which can be verified in each specific application.[20]

At time t, government preferences are given by a function $\Psi^g(k_t, a_t, s_t, a_{t+1}, s_{t+1}, a_{t+2}, s_{t+2}, \ldots)$. We assume that this function is also weakly separable in k_t and its other arguments. For each given sequence of government actions $\{a_s\}_{s=t}^{\infty}$, a unique competitive equilibrium exists. The resulting sequence of private sector actions is given by a sequence $\{s_s\}_{s=t}^{\infty}$, which is independent of k_t, since household preferences are also separable in k_t. We thus specify the government utility from its sequence of actions as that experienced in the competitive equilibrium associated with those actions. With this specification, government preferences can be represented as in equation (8), and an organizational equilibrium can be defined in the same way as in Section II. Existence of an organizational equilibrium is guaranteed by Proposition 1 when Assumptions 2 and 3 hold. However, these assumptions are significantly more restrictive in tax applications. As is well known, optimal tax problems frequently feature nonconvexities, in which case existence may have to be established in the specific application, as we do in our examples. Moreover, anticipation effects from the competitive equilibrium imply

[20]Non-uniqueness can be accommodated by assuming a selection rule on how households coordinate when multiple equilibria are possible, as long as this rule has the properties that the continuation of a selected competitive equilibrium is selected itself as a continuation competitive equilibrium and that the selection is continuous.

that Assumption 3 often does not hold either. It is worth noting that this assumption can be weakened. Its central role in our proof of Proposition 1 is to establish that the continuation sequence $(a^{\bar{E}}_{t+1}, a^{\bar{E}}_{t+2}, \ldots)$ in equation (1) can be made independent of the current deviation a. In our policy applications, we prove this result by showing instead that the static best-response $\arg\max_a V(a, a_0, a_1, a_2, ...)$ is independent of the sequence $\{a_t\}_{t=0}^{\infty}$: hence, any continuation which deters deviation to this action will also be sufficient to deter deviation to any other choice.

As we did for the simpler case of Section III, we relate an organizational equilibrium to a strategic notion of equilibrium. To do so, we need to keep track of histories of play. A symmetric history of play is a record of all actions taken in the past; we distinguish between histories at which the government is called to play, which are given by $h^0 := \emptyset$ and

$$h^t := (a_0, s_0, a_1, s_1, ..., a_{t-1}, s_{t-1}), \quad t > 0,$$

and histories at which households are called to play, that take the form of $h^{p,0} := a_0$ and

$$h^{p,t} := (a_0, s_0, a_1, s_1, ..., a_{t-1}, s_{t-1}, a_t), \quad t > 0.$$

Let H be the set of histories at which the government is called to play, and H^p the set of histories at which households are called to play. For the reasons discussed above, we only keep track of histories in which almost all households have taken the same action.

A strategy for the households is a mapping $\sigma^p : H^p \to S$; likewise, a government strategy is a mapping $\sigma : H \to A$. A *symmetric strategy profile* is a pair (σ^p, σ), representing how all households and the government will act following any symmetric history; it recursively induces a path of play $\{a_t, s_t\}_{t=0}^{\infty}$.

A symmetric strategy profile (σ^p, σ) is a sequential equilibrium if the following is true:

- Given that the government will follow σ and other households will follow σ^p, the actions dictated by σ^p are optimal for each household. After any history $h^{p,t}$, each household takes as given the government policy action a_t and the initial state k_t, which is recursively determined by the history of past play. Moreover, the strategy σ^p followed by other households and the government strategy σ determine the *future* path of aggregate play, $\{s_v, a_{v+1}\}_{v=t}^{\infty}$. Household optimality requires that the sequence of actions prescribed by σ^p is optimal along this path: equivalently stated, it requires the actions prescribed by σ^p to be a competitive equilibrium from period t on, following any arbitrary (symmetric) history.
- Given that households will follow the strategy σ^p and that future governments will follow the strategy σ, and given any past history h^t, the current government choice $\sigma(h^t)$ is optimal.

Proposition G.2. *Given any organizational equilibrium, there exists a sequential equilibrium whose outcome coincides with the organizational equilibrium.*

Proof. Let $(a^*_0, a^*_1, a^*_2, ...)$ be an organizational equilibrium, and let $(s^*_0, s^*_1, ...)$ be the competitive-equilibrium associated with it. We construct a strategy profile recursively as follows:

- $\sigma(\emptyset) = a_0^*$;
- For any $t > 0$ and any history $h^t = (a_0, ..., a_{t-1})$ such that $a_s = a_s^* \quad \forall s = 0, ..., t-1$, $\sigma(h^t) = a_t^*$;
- For any $t > 0$ and any history $h^t = (a_0, ..., a_{t-1})$ such that $\exists s : a_s \neq a_s^*$, define $T := \max\{s < t : a_s \neq \sigma(a_0, ...a_{s-1})\}$ and set $\sigma(h^t) = a_{t-1-T}^*$.
- For any history $h_t^p = (a_0, s_0, a_1, s_1, ..., a_{t-1}, s_{t-1}, a_t)$ at which households are called to play, let $\{a_s^e\}_{s=t+1}^{\infty}$ be the sequence of government actions that follow from period $t+1$ if the government plays the continuation of the strategy σ defined above following $(a_0, ..., a_t)$. Set $\sigma^p(h_t^p)$ to be the competitive equilibrium that is associated with $(a_t, a_{t+1}^e, a_{t+2}^e, ...)$, which exists and is unique by assumption.

By construction, the household strategy satisfies the second condition for a sequential equilibrium for any history of play. For the government, following any history, the strategy prescribes to play the organizational equilibrium sequence, either from its beginning or from some element a_t^*, $t > 0$. Should the government deviate from its strategy, the continuation strategy restarts the organizational equilibrium sequence from a_0^*. By the definition of an organizational equilibrium, continuing along the sequence is always weakly preferred to playing the best one-shot deviation followed by a restart; hence, the government optimality condition is satisfied and the strategy above describes a sequential equilibrium. □

H Details for Section IV

We first provide the details on the separable property of the model environment. Recall that the law of motion of the stock of carbon is given by

$$q_{1t} = q_{1t-1} + \varphi_L e_t,$$
$$q_{2t} = \varphi q_{2t-1} + (1 - \varphi_L)\varphi_0 e_t.$$

It follows that

$$q_{1t} = q_{1,-1} + \varphi_L A \sum_{j=0}^{t} (1 - n_j),$$

$$q_{2t} = \varphi^t q_{2,-1} + (1 - \varphi_L)\varphi_0 A \sum_{j=0}^{t} \varphi^{t-j}(1 - n_j),$$

$$q_t = q_{1t} + q_{2t} = q_{1,-1} + \varphi^t q_{2,-1} + A \sum_{j=0}^{t} (\varphi_L + (1 - \varphi_L)\varphi_0 \varphi^{t-j})(1 - n_j).$$

Given a sequence of $\{s_t\}$ and $\{n_t\}$, it implies that

$$\log c_t = \log(1 - s_t) - \gamma q_t + \alpha \log k_t + (1 - \alpha - \nu) \log n_t + \nu \log(A(1 - n_t))$$
$$= \log(1 - s_t) - \gamma \left(q_{1,-1} + \varphi^t q_{2,-1} + A \sum_{j=0}^{t} (\varphi_L + (1 - \varphi_L)\varphi_0 \varphi^{t-j})(1 - n_j) \right)$$

$$+ \alpha \log k_t + (1-\alpha-\nu)\log n_t + \nu \log(1-n_t) + \nu \log A$$

Turn to the part involving the saving rates. The sequence of capital is

$$\log k_t = \alpha^t \log k_0 + \sum_{j=0}^{t-1} \alpha^{t-j-1} \log s_j$$

The lifetime utility is therefore separable between initial states $(q_{1,-1}, q_{2,-1}, k_0)$ and the sequence of labor and saving rates

$$\begin{aligned} U_0 =& \log c_0 + \delta \sum_{j=1}^{\infty} \beta^j \log c_{t+j} \\ =& G(k_0, q_{1,-1}, q_{2,-1}) + W(s_0, s_1, \ldots) + V(n_0, n_1, \ldots). \end{aligned}$$

Here, $W(s_0, s_1, \ldots)$ captures the impact of saving rates. Using the expression for $\log c_t$ and $\log k_t$, we have

$$W(s_0, s_1, \ldots) = \log(1-s_0) + \frac{\delta\alpha\beta}{1-\alpha\beta}\log(s_0) + \delta \sum_{j=1}^{\infty} \beta^j \left(\log(1-s_j) + \frac{\alpha\beta}{1-\alpha\beta}\log(s_j) \right),$$

which is similar to the baseline quasi-geometric discounting model.

Next, consider the part involving the labor choice.

$$\begin{aligned} & V(n_0, n_1, \ldots) \\ =& (1-\alpha-\nu)\log n_0 + \nu\log(1-n_0) - \gamma A \left(\varphi_L \frac{1-\beta+\delta\beta}{1-\beta} + (1-\varphi_L)\varphi_0 \frac{1-\varphi\beta+\delta\varphi\beta}{1-\varphi\beta} \right)(1-n_0) \\ & + \delta\beta \left(-\gamma A \left(\frac{\varphi_L}{1-\beta} + \frac{\varphi_0(1-\varphi_L)}{1-\beta\varphi} \right) \sum_{j=0}^{\infty} \beta^j (1-n_{j+1}) + (1-\alpha-\nu) \sum_{j=0}^{\infty} \beta^j \log n_{j+1} + \nu \sum_{j=0}^{\infty} \beta^j \log(1-n_{j+1}) \right) \\ =& (1-\alpha-\nu)\log n_0 + \nu\log(1-n_0) - \gamma A \left(\varphi_L \frac{1-\beta+\delta\beta}{1-\beta} + (1-\varphi_L)\varphi_0 \frac{1-\varphi\beta+\delta\varphi\beta}{1-\varphi\beta} \right)(1-n_0) \\ & - \beta(1-\delta)\Big((1-\alpha-\nu)\log n_1 + \nu\log(1-n_1) - \gamma A (\varphi_L + (1-\varphi_L)\varphi_0)(1-n_1) \Big) \\ & + \beta V(n_1, n_2, \ldots). \end{aligned}$$

Characterization of steady state The organizational equilibrium requires that $V(n_0, n_1, \ldots) = V(n_1, n_2, \ldots) = \overline{V}$, which leads to

$$\begin{aligned} & -\gamma A \left(\varphi_L \frac{1-\beta+\delta\beta}{1-\beta} + (1-\varphi_L)\varphi_0 \frac{1-\varphi\beta+\delta\varphi\beta}{1-\varphi\beta} \right)(1-n_0) + (1-\alpha-\nu)\log n_0 + \nu\log(1-n_0) \\ & = \beta(1-\delta)\Big((1-\alpha-\nu)\log n_1 + \nu\log(1-n_1) - \gamma A (\varphi_L + (1-\varphi_L)\varphi_0)(1-n_1) \Big) + (1-\beta)\overline{V}. \end{aligned}$$

In the steady state, the best constant action maximizes the following object

$$n^O = \underset{n}{\text{argmax}} -\gamma A \left(\varphi_L \left(\frac{1-\beta+\delta\beta}{1-\beta} - \beta(1-\delta) \right) + (1-\varphi_L)\varphi_0 \left(\frac{1-\varphi\beta+\delta\varphi\beta}{1-\varphi\beta} - \beta(1-\delta) \right) \right)(1-n)$$
$$+ (1-\alpha-\nu)(1-\beta(1-\delta))\log n + \nu(1-\beta(1-\delta))\log(1-n).$$

The first-order condition implies

$$\Lambda^O + (1-\alpha-\nu)\frac{1}{A}\frac{1}{n^O} = \nu\frac{1}{A}\frac{1}{1-n^O},$$

where Λ^O in the organizational equilibrium is given by

$$\Lambda^O \equiv \gamma \left(\varphi_L \left(1 + \frac{\delta\beta}{(1-\beta)(1-\beta(1-\delta))} \right) + (1-\varphi_L)\varphi_0 \left(1 + \frac{\delta\varphi\beta}{(1-\varphi\beta)(1-\beta(1-\delta))} \right) \right).$$

When $\delta = 1$, the steady-state policy reconciles with the outcome characterized in Golosov et al. (2014), which corresponds to the Ramsey outcome in the long-run

$$\Lambda^R \equiv \gamma \left(\varphi_L \frac{1}{1-\beta} + (1-\varphi_L)\varphi_0 \frac{1}{1-\varphi\beta} \right).$$

The allocation of labor in the Markov equilibrium solves the action payoff taken future labor choice as given

$$n^M = \text{argmax} -\gamma A \left(\varphi_L \frac{1-\beta+\delta\beta}{1-\beta} + (1-\varphi_L)\varphi_0 \frac{1-\varphi\beta+\delta\varphi\beta}{1-\varphi\beta} \right)(1-n) + (1-\alpha-\nu)\log n + \nu\log(1-n).$$

The implied tax is

$$\Lambda^M = \gamma \left(\varphi_L \frac{1-\beta+\delta\beta}{1-\beta} + (1-\varphi_L)\varphi_0 \frac{1-\varphi\beta+\delta\varphi\beta}{1-\varphi\beta} \right).$$

I Details for Section V

Static equilibrium Recall that the final goods is an aggregator of the two intermediate goods

$$y_t = \left[0.5^{1-\rho} m_{1t}^{\rho} + 0.5^{1-\rho} m_{2t}^{\rho} \right]^{\frac{\rho-1}{\rho}}.$$

The aggregate price index $\mathcal{P}_t$ is given by

$$\mathcal{P}_t = \left[0.5 p_{1t}^{\frac{\rho}{\rho-1}} + 0.5 \right]^{\frac{\rho-1}{\rho}}$$

. The demand schedules for goods 1 and 2 satisfy

$$m_{1t} = \frac{1}{2}\left(\frac{p_{1t}}{\mathcal{P}_t}\right)^{\frac{1}{\rho-1}} y_t, \qquad m_{2t} = \frac{1}{2}\left(\frac{1}{\mathcal{P}_t}\right)^{\frac{1}{\rho-1}} y_t,$$

The production functions in the two sectors are given by

$$y_{1t} = AL_{1t}^{1-\alpha}k_{1t}^{\alpha}k_t^{1-\alpha}, \quad \text{and} \quad y_{1t} = L_{2t}^{1-\alpha}k_{1t}^{\alpha}k_t^{1-\alpha}.$$

We impose fixed labor input in the two sectors, i.e., $L_{1t} = L_{2t} = 1$.

Denote the tariff rate as τ_t. In the home country, the price of goods 2 is 1 by normalization. Therefore, the price of goods 2 at the foreign country is $\frac{1}{1+\tau_t}$ due to the law of one price and the symmetric tariff rate assumption. Again, by symmetry, the price of goods 1 in country 1 is also $p_{1t} = \frac{1}{1+\tau_t}$. That is, a higher tariff rate lowers the more productive sector's relative price. It also implies that there is a one-to-one mapping between the tariff rate τ_t and the goods 1 price p_{1t}.

Since capital is free to flow across sectors, the return to capital is equalized across the two sectors

$$p_{1t}Ak_{1t}^{\alpha-1}k_t^{1-\alpha} = k_{2t}^{\alpha-1}k_t^{1-\alpha}$$

which leads to

$$k_{2t} = \phi_t k_t \qquad \text{where} \qquad \phi_t = \frac{1}{1+(p_{1t}A)^{\frac{1}{1-\alpha}}}.$$

It also follows that the output in the two sectors are given by

$$y_{1t} = A(1-\phi_t)^{\alpha}k_t, \quad y_{2t} = \phi_t^{\alpha}k_t,$$

and the nominal GDP in country 1 can be expressed as

$$v_t = p_{1t}Ak_{1t}^{\alpha}k_t^{1-\alpha} + k_{2t}^{\alpha}k_t^{1-\alpha} = k_t(p_{1t}A(1-\phi_t)^{\alpha} + \phi_t^{\alpha})) = k_t\phi_t^{\alpha-1}.$$

The total expenditure and the nominal GDP is identical, which implies that

$$v_t = \mathcal{P}_t y_t = p_{1t}m_{1t} + m_{2t},$$

where the total usage of goods 1 and 2 follows from the demand schedule

$$m_{1t} = \frac{1}{2}\left(\frac{p_{1t}}{\mathcal{P}_t}\right)^{\frac{1}{\rho-1}} \frac{v_t}{\mathcal{P}_t}, \quad m_{2t} = \frac{1}{2}\left(\frac{1}{\mathcal{P}_t}\right)^{\frac{1}{\rho-1}} \frac{v_t}{\mathcal{P}_t}.$$

By symmetry, the terms of the trade is 1 in equilibrium. Denote the tariff revenue as $\mathcal{R}$, which can be derived as

$$\mathcal{R}_t = (m_{2t} - y_{2t})\left(\frac{1}{p_{1t}} - 1\right) = k_t\phi_t^{\alpha}\left(\frac{1}{\phi_t}\frac{1}{2}\mathcal{P}_t^{-\frac{\rho}{\rho-1}} - 1\right)\left(\frac{1}{p_{1t}} - 1\right).$$

Under the assumption that the tariff revenue is equally split between the two groups of workers, the consumption of workers in industry 1 and 2 are

$$c_{1t} = \frac{\frac{1}{2}\mathcal{R}_t + w_{1t}}{\mathcal{P}_t} = \frac{\frac{1}{2}\mathcal{R}_t + (1-\alpha)p_{1t}A(1-\phi_t)^{\alpha}k_t}{\mathcal{P}_t} = \chi_1(\tau_t)k_t, \tag{18}$$

$$c_{2t} = \frac{\frac{1}{2}\mathcal{R}_t + w_{2t}}{\mathcal{P}_t} = \frac{\frac{1}{2}\mathcal{R}_t + (1-\alpha)\phi_t^{\alpha}k_t}{\mathcal{P}_t} = \chi_2(\tau_t)k_t, \tag{19}$$

where p_{1t}, $\mathcal{P}_t$, ϕ_t are all functions of τ_t.

Capitalists The real return to capital is

$$r(\tau_t) = \frac{\alpha k_{2t}^{\alpha-1}k_t^{1-\alpha}}{\mathcal{P}_t} = \frac{\alpha\phi_t^{\alpha-1}}{\mathcal{P}_t}.$$

The problem of capitalists can be written as

$$\max \sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\sigma}}{1-\sigma}$$

subject to

$$c_t + k_{t+1} = (r(\tau_t) + 1 - \delta)k_t.$$

The Euler equation is

$$c_t^{-\sigma} = \beta(r(\tau_{t+1}) + 1 - \delta)c_{t+1}^{-\sigma}.$$

Denote s_t as the saving rate, the Euler equation can be expressed as

$$\left(\frac{1-s_t}{s_t}\right)^{-\sigma} = \beta\left(r(\tau_{t+1}) + 1 - \delta\right)^{1-\sigma}\left(1 - s_{t+1}\right)^{-\sigma}.$$

Meanwhile, given a sequence of saving rates and tariff rates, the capital evolution satisfies

$$k_t = \prod_{j=0}^{t-1} s_j\left(r(\tau_j) - \delta\right)k_0.$$

Welfare Suppose the policymaker's preference is to maximize

$$U = \sum_{t=0}^{\infty} \beta^t \Big(\lambda \log c_{1t} + (1-\lambda)\log c_{2t}\Big)$$

Define

$$\chi(\tau) = \lambda \log \chi_1(\tau) + (1-\lambda)\log \chi_2(\tau),$$

where $\chi_1(\cdot)$ and $\chi_2(\cdot)$ are defined in equation (18) and (19).

The total welfare is separable between capital and the trade policy

$$U = \frac{1}{1-\beta}\log k_0 + \sum_{t=0}^{\infty}\beta^t \chi(\tau_t) + \frac{\beta}{1-\beta}\sum_{t=0}^{\infty}\beta^t(\log s_t + \log(r(\tau_t)+1-\delta)).$$

Markov Equilibrium The policymaker in the Markov equilibrium takes future tariff rates as given and simply maximizes the following object

$$\chi(\tau) + \frac{\beta}{1-\beta}\log(r(\tau)+1-\delta)$$

which yields a constant tariff. Note that the impact of tariff on the saving rate is not taken into account.

Organization Equilibrium The steady-state allocation in the organizational equilibrium satisfies

$$\max_{s,\tau} \chi(\tau) + \frac{\beta}{1-\beta}\log(r(\tau)+1-\delta) + \frac{\beta}{1-\beta}\log s$$

subject to

$$\left(\frac{1-s}{s}\right)^{-\sigma} = \beta\left(r(\tau)+1-\delta\right)^{1-\sigma}(1-s)^{-\sigma}$$

With $\sigma < 1$, the saving rate is decreasing in τ, which gives the policymaker a larger incentive to lower the tariff rate.

Ramsey Outcome With commitment, the policymaker's problem is

$$\max_{\{\tau_0,\tau_1,\ldots\}} \sum_{t=0}^{\infty}\beta^t\chi(\tau_t) + \frac{\beta}{1-\beta}\sum_{t=0}^{\infty}\beta^t(\log s_t + \log(r(\tau_t)+1-\delta))$$

subject to

$$\left(\frac{1-s_t}{s_t}\right)^{-\sigma} = \beta\left(r(\tau_{t+1})+1-\delta\right)^{1-\sigma}(1-s_{t+1})^{-\sigma}$$

Let $\beta^t\mu_t$ denote the multiplier associated with the constraint involving s_t and s_{t+1}. For τ_0, the choice is to maximize

$$\chi(\tau_0) + \frac{\beta}{1-\beta}\log(r(\tau_0)+1-\delta)$$

For s_0, the first-order condition is

$$\frac{\beta}{1-\beta}\frac{1}{s_0} = -\mu_0\sigma\left(\frac{1-s_0}{s_0}\right)^{-\sigma-1}\frac{1}{s_0^2}$$

For $t \geq 1$, the first-order condition with respect to τ_t is

$$\beta^t\left(\chi_\tau(\tau_t) + \frac{\beta}{1-\beta}\frac{r_\tau(\tau_t)}{r(\tau_t)+1-\delta}\right) = \beta^{t-1}\mu_{t-1}(1-\sigma)\beta\left(r(\tau_t)+1-\delta\right)^{-\sigma}(1-s_t)^{-\sigma}r_\tau(\tau_t).$$

The first order condition with respect to s_t is

$$\beta^t \left(\frac{\beta}{1-\beta} \frac{1}{s_t} \right) = \beta^{t-1} \mu_{t-1} \sigma \beta \left(r(\tau_t) + 1 - \delta \right)^{1-\sigma} (1-s_t)^{-\sigma-1} - \beta^t \mu_t \sigma \left(\frac{1-s_t}{s_t} \right)^{-\sigma-1} \frac{1}{s_t^2}.$$

These conditions characterize the transition dynamics.

References

Abreu, Dilip, David Pearce, and Ennio Stacchetti, "Optimal Cartel Equilibria with Imperfect Monitoring," *Journal of Economic Theory*, 1986, *39*, 251–269.

_ , _ , and _ , "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 1990, *58* (5), 1041–1063.

Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski, "Optimal taxes on fossil fuel in general equilibrium," *Econometrica*, 2014, *82* (1), 41–88.

Jackson, Matthew O. and Leeat Yariv, "Present Bias and Collective Dynamic Choice in the Lab," *American Economic Review*, 2014, *104* (12), 4184–4204.

_ and _ , "Collective Dynamic Choice: The Necessity of Time Inconsistency," *American Economic Journal: Microeconomics*, 2015, *7* (4), 150–78.

Jovanovic, Boyan and Robert W. Rosenthal, "Anonymous Sequential Games," *Journal of Mathematical Economics*, 1988, *17* (1), 77–87.

Kocherlakota, Narayana R., "Reconsideration-Proofness: A Refinement for Infinite Horizon Time Inconsistency," *Games and Economic Behavior*, July 1996, *15* (1), 33–54.