

*On-line Appendix for
Rules With Discretion and Local Information*

*Renee Bowen, David M. Kreps, and Andrzej Skrzypacz
Graduate School of Business, Stanford University
March 2013*

This appendix to “Rules With Discretion and Local Information” collects a number of extensions and variations on the main model and results of that paper. It assumes that the reader has first read the main paper.

Random alarm clocks

In a number of the constructions of this appendix, we rely on so-called *random alarm clocks*, which in most cases are used to set the duration of “punishment” for a player who has failed to do a favor. These random alarm clocks always have durations that are exponentially distributed, meaning that if the clock is “set” at some time T_0 , it goes off at time $T_0 + t$ or earlier with probability $1 - e^{-\phi t}$, where ϕ is the parameter that describes the distribution. These alarm clocks are in all cases meant to be publicly observable, with the further property that until the clock “rings,” no one knows anything more than that it has yet to ring. Of course, by the memoryless property of the exponential distribution, this means that a clock set at time T_0 with parameter ϕ that has not gone off by time $T_1 > T_0$ will now go off between time T_1 and $T_1 + t$ with probability $1 - e^{-\phi t}$.

Random alarm clocks with these characteristics are fairly straightforward generalizations of publicly observable randomizations, which in turn are a staple of analysis in these sorts of models. Indeed, a random alarm clock with parameter ϕ is easily approximated to any desired degree by a sequence of standard publicly observable randomizations: Fixing a starting time T_0 and some small $\delta > 0$, at dates $T + k\delta$ for $k = 1, \dots$, a publicly observable randomization is conducted with outcomes “ring” having probability $\phi\delta$ and “not yet” with probability $1 - \phi\delta$. The first occurrence of the outcome “ring” corresponds to the clock ringing (at which point, of course, the sequence of public randomizations can stop). As δ approaches zero, this approaches the random alarm clock described in the previous paragraph in all respects.

No cheap talk, but unilateral discretion

Suppose that, within the context of our basic model, we give favor receivers no cheap-talk role at all. When an i -for- j favor opportunity arises, the favor giver i observes the cost-benefit vector (x, y) and decides whether to do the favor. The favor receiver j sees these same things; all other players observe that an i -for- j favor opportunity has arisen and whether i does the favor or not.

Following the spirit of the paper, where we are looking for equilibria that involve social enforcement, we look for perfect social equilibria for this game.

Consider a selection \mathcal{S} such that, for each i and j , \mathcal{S}_{ij} is either the full support of the random vector (x_{ij}, y_{ij}) or the null set. Suppose that, for all i and j such that \mathcal{S}_{ij} is the full support of (x_{ij}, y_{ij}) ,

$$\max \{x : (x, y) \text{ is in the support of } (x_{ij}, y_{ij}) \text{ for some } y\} \leq \frac{\mathcal{A}_i(\mathcal{S}) - \mathcal{B}_i(\mathcal{S})}{r}.$$

Then the selection \mathcal{S} can be implemented as a perfect social equilibrium in obvious fashion: When an i -for- j favor opportunity arises, if this ordered pair is such that \mathcal{S} calls for i to do all favors for j , i does the favor. And if she fails to do so, everyone moves to autarky.

The more interesting cases are where, for ordered pairs (i, j) , the selection \mathcal{S} calls for i to do some favors for j , but not all possible favors. It is those cases that we address here.

We reiterate a point made in the paper: With three or more players, for such selections \mathcal{S} , one cannot implement the selection \mathcal{S} in the sense of the paper; that is, where along the path of play, all favors in \mathcal{S} and only favors in \mathcal{S} are done. If (i, j) are an ordered pair for which \mathcal{S} specifies some but not all favors, i 's continuation value (when an i -for- j favor arises) if she does the favor must strictly exceed her continuation value if she fails to do the favor; otherwise, she would do no favors. Moreover, there must be cases along the path of equilibrium play in which she does favors, and others where she does not. Hence both continuation values must be reached along the path of equilibrium play, which is inconsistent with (time-homogeneous) implementation of \mathcal{S} .

Still, we can imagine perfect social equilibria in which i , at least at her first opportunity to do favors for j , exercises unilateral discretion in which favors she does, and is punished if she fails to do some. For instance, and in the spirit of the equilibrium of Proposition 1 in the paper, we can imagine that, if i fails to do an i -for- j favor, a public randomization occurs, in which with some probability ϕ_{ij} , the failure to do the favor is ignored; with complementary probability, all players move forever to autarky. In this case, i would choose to do the i -for- j favor (x, y) if and only if

$$v_i - x \geq \phi_{ij}v_i, \tag{OA.1}$$

where v_i is i 's expected payoff if all players behave in like fashion. Note that finding vector of values $(v_i)_{i=1, \dots, I}$ becomes a fixed-point problem; in this scenario, unlike in the scenario of Proposition 1 in the paper, undone favors lie along the equilibrium path, so unless all favors pass the test implicit in (OA.1) when $v_i = \sum_{j \neq i} [\lambda_{ji}a_{ji} - \lambda_{ij}b_{ij}]/r$, (in which case the all-favors selection is consistent with an equilibrium for the given probabilities ϕ_{ij}), the play degenerates into autarky (almost surely) in finite time. Note that the fixed-point problem is non-trivial, since excluding some favors may accelerate the time at which play moves to autarky but, at the same time, it may increase the expected flow of net benefits to players, especially if favors that are ruled out provide little benefit at great cost.

This is but one form of (possible) equilibrium with unilateral discretion, based on one type of punishment. A different sort of equilibrium somewhat resembles the equilibrium of Proposition 2, in that

each player, at any point of the game, is in one of two states, G (for grace) or P (purgatory). Imagine that each player i , if she is in state G, does a selection of favors with an overall arrival rate λ and expected cost (conditional on arrival of one of these favors) c . Favors that she refuses to do arrive at rate λ' . When she refuses to do a favor, she moves to state P, in which she must do all favors; if any player in state P fails to do any favor, all players move immediately to autarky, forever. Favors of the second type have an expected cost (conditional on being such a favor) of c' . Finally, the duration of her stay in P is governed by an exponential alarm clock with parameter ϕ (which may be 0, meaning she never leaves state P); when she leaves P, she returns to G. All the random events described are independent of one another.

Before continuing with the analysis, note that we have departed in several respects from the construction in Proposition 2. We retain the two-state (G and P) structure used there. But in Proposition 2, state P involved the punishment of receiving no favors, while continuing to do all favors called for in the selection \mathcal{S} . Here, state-P punishment consists of having to do favors that one doesn't have to do while in state G; however favors continue to be done for someone in P. And whereas a player in state P in Proposition 2 could get back into state G by doing favors (and “passing” a randomized screen imposed by the favor receiver), here the duration of punishment is out of the control of the player in state P; she does all favors until she is released by the random alarm clock.

For such a regime to constitute a perfect social equilibrium, two sorts of incentive compatibility constraints need checking. Players in state P are required to do all favors: The highest-cost favor that a player is ever called upon to do must be less than her continuation value, starting in state P. (Recall that autarky has continuation value 0.) And the difference in continuation values in states G and P must be such that each player will do the “right” selection of favors when in state G, no more and no less.

Please note that, in this sort of regime, a player has no influence whatsoever on which favors are done for her. Fellow players may (in general, will) drift in and out of state P, and while they are in state P, other players may get more favors done for them. For this reason, it may be difficult to compute the precise continuation value of being in state P, but we can compute a lower bound on this—which may be sufficient to check to first sort of incentive compatibility constraint—if we assume that the flow of benefits is what the player gets when everyone else is in state G. On the other hand, we need for the second incentive compatibility constraint to know the difference between the continuation values in G and in P; since each player has no influence (and her state has no influence) on the flow of benefits she receives, we can compute the difference in continuation values by compute the difference in her expected costs.

Let C denote her discounted costs starting in G, and let C' be the expectation starting in P. Conditioning on the first arrival of a favor to be done in state G, we have

$$\begin{aligned} C &= \int_0^\infty e^{-rt} \left[\frac{\lambda}{\lambda + \lambda'} (c + C) + \frac{\lambda'}{\lambda + \lambda'} C' \right] (\lambda + \lambda') e^{-(\lambda + \lambda')t} dt \\ &= \frac{\lambda}{\lambda + \lambda' + r} (c + C) + \frac{\lambda'}{\lambda + \lambda' + r} C' \end{aligned}$$

For C' , we condition on the earliest of: Arrival of a favor of one type or the other; ringing of the alarm clock:

$$C' = \int_0^{\infty} e^{-rt} \left[\frac{\lambda(c + C') + \lambda'(c' + C') + \phi C}{\lambda + \lambda' + \phi} \right] (\lambda + \lambda' + \phi) e^{-(\lambda + \lambda' + \phi)t} dt$$

$$= \left[\frac{\lambda(c + C') + \lambda'(c' + C') + \phi C}{\lambda + \lambda' + \phi + r} \right].$$

This gives us two linear equations in two unknowns to solve.

In fact, the important quantities from the perspective of equilibrium verification are C' and $C' - C$. These are

$$C' = \frac{\lambda' c'}{r} \cdot \frac{\lambda' + r}{\phi + r + \lambda'} + \frac{\lambda c}{r} \quad \text{and} \quad C' - C = \frac{\lambda' c'}{\phi + r + \lambda'}.$$

Note also that

$$C = \frac{\lambda' c'}{r} \cdot \frac{\lambda'}{\phi + r + \lambda'} + \frac{\lambda c}{r}.$$

The last is important only to note that her expected costs are decreasing in ϕ ; this is intuitively obvious, as larger ϕ means less time doing both sorts of favors.

To see how these computations work in a specific example, consider the circle example from the paper, where each i -for- $i + 1$ favor has cost-benefit vector with the following distribution: $(x, y) = (2, 3)$ with probability 0.8; $(x, y) = (1, 4)$ with probability 0.1; and $(x, y) = (5, 4)$ with probability 0.1. We would like to create an equilibrium in which each player, while in state G, does only the first two types of favor; the third is excluded because its cost exceeds its benefit. We take $r = 0.1$

In terms of the notation above, $\lambda = 0.9$, $\lambda' = 0.1$, $c = 17/9$, and $c' = 5$. The largest favor that the player is expected to do in state P has cost 5. A lower bound on the expected flow rate of benefits is the expected flow rate when all other players are in state G (or, for the circle, when her left-hand neighbor is in state G), which is 2.8 units per unit time, for an expected net present value of 28. At the moment a player is called upon to do a favor in state P, she must incur that immediate cost, which is 5 or less, and she has an expected value of ongoing costs equal to C' for these data, which is largest when $\phi = 0$, at which value $C' = 22$. (If $\phi = 0$, she stays in P forever; there is no escape. In that case, the flow rate of favors she must do is 2.2 per unit time, hence $C' = 22$.) So, with an alternative of triggering autarky, she is willing to do any favor with immediate cost 6 or less.

The final equilibrium check is: Will she do favors of cost 2 or less when in state G, if not doing them sends her to state P? Since her decision has no impact on her expected discounted benefits from favors

done for her, it is a matter of computing $C' - C$, or how much more it will cost her (in expectation) to be in P instead of in G. From the computations above, this is

$$\frac{\lambda'c'}{\phi + r + \lambda'} = \frac{0.5}{\phi + 0.2};$$

this must equal or exceed 2 (so that she does favors of cost 2 or more), which is

$$\frac{0.5}{\phi + 0.2} \geq 2 \quad \text{or} \quad \phi \leq 0.05.$$

(Note that even if $\phi = 0$, which maximizes $C' - C$, the difference is 2.5, so she will not do favors of cost 5 to avoid moving to state P. This is, in fact, obvious: The difference between P and G is that favors of cost 5 must be done in state P, and doing such a favor today to avoid having to do it “tomorrow” cannot be optimal when benefits and costs are discounted.)

To reiterate, “punishment” in this sort of regime for having failed to do a favor amounts to requiring the miscreant to do all favors for some length of time. A more severe punishment—which might let us improve on the all-favors equilibrium even further, or provide an equilibrium in cases where r is too large for the mild-punishment equilibrium to work—would be to *withhold* favors from a player, at least for a while. Imagine, for instance, that any player in state P is required to do all favors she is called upon to do (except for other players who are also in state P), but no one is required to do any favors for her. (If she fails to do any favor, play moves to autarky.) For something like this to be an equilibrium, we must restore a player being punished to state G in finite (expected) time; if we didn’t, she wouldn’t have any incentive to do favors for others. In other words, the condition that this punishment is not so severe that continuation values when in state P are less than the most expensive favor the player may be called upon to do may bind. But, beyond this, the possibility of such more-onerous punishments seems alluring.

Alluring or not, computing values for equilibria of this sort is more complex than for the mild-punishment equilibria we have constructed. The problem is that while the full “system” of states of the players forms a time-homogeneous, continuous-time Markov chain, the status of any single player does not. Suppose, for simplicity, that players are arranged in a circle, where each player (only) does favors for her clockwise neighbor. Suppose that player i is in state G. How long she stays there depends on the state of her clockwise neighbor: If he is also in state G, then she must either perform favors for him or be moved into state P; if he is in state P, she is safe, at least for a while. In theory, one can set up a system of simultaneous linear equations to determine equilibrium values. But, in general, the number of equations and unknowns is on the order of $I \times 2^I$.¹ Because of the complexities involved, we have not pursued the computation of equilibria with unilateral discretion and this sort of punishment. (But see later on in this appendix, where we look at other equilibria with ex-post cheap talk.)

¹ This represents the number of players I times the number of states of the grand Markov process, although 2^I is an overstatement, as the state can never have all players in state p simultaneously. Of course, in symmetric games and for symmetric equilibria, we can compute for one player only.

While we know that, with unilateral discretion, we cannot implement the time-homogeneous application of a selection \mathcal{S} , perhaps we can implement efficient outcomes with non-stationary equilibria. The idea, roughly, would be that if i fails to do a favor for j , j receives compensation in the continuation play. At least for the u-efficient outcome, this is impossible; the only way for the u-efficient payoffs to be realized is for the u-efficient and only the u-efficient favors to be done.

But, in the spirit of the Folk Theorem, perhaps it is possible that, for a given selection \mathcal{S} , we can find perfect social equilibria for small interest rates r that, in terms of the normalized payoffs, approach the payoffs that would be obtained from time-homogeneous application of \mathcal{S} . Under certain conditions, this is possible:

Proposition OAI. *Suppose that selection \mathcal{S} satisfies*

$$\mathcal{A}_i(\mathcal{S}) > \mathcal{B}_i(\mathcal{S}) \quad \text{for all } i, \text{ and} \quad (\text{A}')$$

$$\text{if } (x, y) \in \mathcal{S}_{ij} \text{ and } (x, y) \notin \mathcal{S}_{ij}, \text{ then } x < x'. \quad (\text{B}')$$

Then for any sequence $\{r_n\}$ of interest rates with $\lim_n r_n = 0$, one can find perfect social equilibria, one for each r_n , (for the game with no cheap talk) such that, if $v_i(n)$ is the expected payoff to player i in the n th of these equilibria,

$$\lim_n r_n v_i(n) = \mathcal{A}_i(\mathcal{S}) - \mathcal{B}_i(\mathcal{S}).$$

A comparison of A' and B' with A and B from Proposition 1 is in order: A' ensures A for small enough r , while B' is nearly the same as B, only with a strict inequality. So, one way to read this proposition is: For small interest rates, ex-ante cheap talk doesn't add any possible (normalized) payoffs to unilateral discretion, at least insofar as B is replaced by B'.

Sketch of the proof. This proposition is proved by more or less standard techniques that are employed in the general treatment of Folk Theorems, albeit those techniques must be adapted to this continuous-time, asynchronous-event setting. So we only provide a sketch.

The method of proof is, essentially, to construct equilibria that, in the limit, as $r \rightarrow 0$, realize the selection- \mathcal{S} (normalized) payoffs. There are a number of ways to do this; we are choosing the one we find easiest to adapt to this context and to understand. The equilibria have the following form:

A very large time period T is chosen. Over this time period, the number of i -for- j favors expected to occur is T/λ_{ij} . Let $\psi_{ij} := \text{Prob}[(x, y) \in \mathcal{S}_{ij}]$, where (x, y) is distributed according to the probability distribution that governs draws of cost-benefit vectors for i -for- j favors. Therefore, the number of i -for- j favors from \mathcal{S}_{ij} expected over the time period is $\psi_{ij}T/\lambda_{ij}$. Let n_{ij} be the integer part of $\psi_{ij}T/\lambda_{ij}$. Then over each time period of the form $[kT, (k+1)T)$, each i is given a quota of n_{ij} favors she can

fail to do for j , for each j . If at any time, any player exceeds her quota *vis a vis* another player j , all players immediately move to autarky. If no player violates any of her $I - 1$ quotas over one of these time intervals, a fresh set of quotas (of the same size for each ordered pair) is put in place for the next time interval of length T .

Note, then, that if we can be assured that no player will ever willingly violate any of her quotas, each player, over each time period of length T , must solve $I - 1$ optimization problems, namely, how to best use up the quota of no-favor opportunities *vis a vis* each (other) player j . Each of these is a complex dynamic-programming problem, but none of these problems for one ordered pair i and j interacts with the solution of a problem for some other pair i' and j' . In particular, as long as we can be assured that no player will ever willingly violate any of her quotas, no player is actively concerned about which favors are done for her; she is only actively concerned with minimizing the cost of favors that she does.

Recall that $M_{ij}(\mathcal{S})$ is the cost of the most expensive i -for- j favor that is in \mathcal{S}_{ij} . Note that condition B' says that $(x, y) \in \mathcal{S}_{ij}$ if and only if $x \leq M_{ij}(\mathcal{S})$.

We assert that, if T is very large and if r is very, very close to zero, so that e^{-rT} is close to one, the solution to each of these cost-minimization problems leads with probability close to one to: Don't do favors whose cost exceeds $M_{ij}(\mathcal{S})$. That is, in each ordered pair i and j , i will omit precisely the favors for j that are not in \mathcal{S}_{ij} . Of course, she can't accomplish this perfectly, since she doesn't know how many i -for- j favor opportunities there will be, when they will occur (some discounting is going on), or what will be their costs. But if T is large, the strong law of large numbers can be employed so that, "on average, and except for an event of small probability," she can and will omit those favors that are not in \mathcal{S} and do those that are in \mathcal{S} .

This means that the expected overall payoff for each players over each of these time intervals is strictly positive; this is where A' comes in. And then, since even for large T , if r is very, very small, the results of any single period of length T loom small in terms of the overall expected payoff, each player prefers to "get through" the current period of length T , no matter how badly things are going (since, in expectation, what is left has bounded cost, even if all favors for the rest of this period must be done because the player has exhausted all her quotas), so she can get to future periods, rather than exceeding a quota and going forevermore to autarky.

Of course, the equilibrium we've outlined is both perfect and social. ■

A partial converse to this result is the following:

Proposition OA2. *If there are more than two players and if, for some pair i and j such that $\lambda_{ij} > 0$, there are favor types $(x(1), y(1))$ and $(x(2), y(2))$ such that $y(1) > x(1) > x(2) > y(2)$. There exists a strictly positive number K such that, if (v_1, \dots, v_I) is the vector of expected payoffs to the respective players in some perfect social equilibrium for the game with interest rate r , then*

$$r[v_1 + \dots + v_I] \leq \sum_i [\mathcal{A}_i(\mathcal{S}^U) - \mathcal{B}_i(\mathcal{S}^U)] - K.$$

In words, u-efficient payoff vector cannot be approached (in terms of normalized payoffs) in any perfect social equilibrium, if for some pair i and j , there is a u-efficient favor the cost of which is strictly greater than the cost of a u-inefficient favor.

Compare with Proposition 3 in the text. This is a similar “anti-folk theorem,” but here there is no cheap talk, and so the condition required to bound perfect-social-equilibrium payoffs away from the u-efficient payoffs is simpler. We omit the proof, since it is a simpler version of the proof of Proposition 3, which is provided in the published Appendix, but to sketch: At each occurrence of an i -for- j favor in any perfect social equilibrium, either i doesn’t do the favor (with probability 1) if it is of type $(x(1), y(1))$, or she does it with probability 1 if it is of type $(x(2), y(2))$. (The continuation values if she does or doesn’t do the favor are the same, and the second type of favor is less costly.) Hence every time an i -for- j favor occurs, there will be a loss in value from u-efficiency.

Alternative equilibria for ex-post cheap talk

In the paper, and specifically in the proof of Proposition 2, we construct an equilibrium with ex-post cheap talk that implements selection \mathcal{S} whenever the obvious incentive-compatibility constraint holds; viz., for all i ,

$$\max_{j \neq i} M_{ij}(\mathcal{S}) \leq \frac{A_i(\mathcal{S}) - B_i(\mathcal{S})}{r}. \quad (\text{A})$$

Since Condition A is clearly necessary for the implementation of \mathcal{S} as a perfect equilibrium (social or not), we can’t do better with other forms of equilibrium in the sense of implementing more selections.

But, as discussed in the paper, the equilibrium that we construct is filled with situations in which players are indifferent between the actions the equilibrium says they take and others that, in equilibrium, they don’t take. And, in some off-path positions, the equilibrium requires favor receivers to employ a mixed strategy concerning his cheap-talk pronouncement.

As discussed in the paper, there is one moment in the game at which this sort of indifference is endemic, if ex-post cheap talk is to be used effectively to distinguish between situations where the favor giver should be forgiven for not doing a favor (because the favor was not called for in the selection \mathcal{S}) and where the favor giver should be punished (because she didn’t do a favor that is called for in \mathcal{S}). Since the favor decision has been made, all that matters to the favor receiver is how his continuation payoffs compare, and in a social equilibrium, these can’t distinguish between the two cases (as long as there are at least three players). So those continuation values must be equal.

But, if the inequalities in (A) contain a lot of slack, and if the most expensive favors players could ever be called upon to do are not too expensive (in a sense to be made precise), alternative equilibria employing ex-post cheap talk can be constructed that implement some selections \mathcal{S} , which involve strict-best responses in all situations except for the one in the previous paragraph.

The key, of course, is in the nature of “punishment” for a favor giver who fails to do a favor and is not excused. In the equilibria of Proposition 2, a malfeasing favor giver is punished by being required to do all favors called for under \mathcal{S} while receiving no favors, and the malfeator is restored to the state of grace (in which the flow of benefits begins again) by (a) doing required favors and then (b) having the favor receiver (with the appropriate probability) restore her to grace. Imagine instead that: (a) Punishment involves being made to do all favors, including those not called for by \mathcal{S} and, perhaps, losing the benefit of favors done for the player being punished, where a player being punished who fails to do any favor triggers autarky for all. And (b) the duration of punishment is out of the control of the players, determined instead by a random alarm clock. We give two specifications that follow these lines, depending on whether a player being punished receives the benefit of any favors.

Fix a selection \mathcal{S} . Let \mathcal{C}_i be the expected flow rate of costs to i , if she is doing all favors; that is,

$$\mathcal{C}_i = \sum_{j \neq i} \lambda_{ij} \mathbf{E}[x_{ij}],$$

where \mathbf{E} denotes expectation of the random cost x_{ij} of each favor. The case where the selection \mathcal{S} calls for some player(s) i to do all favors takes special handling, so temporarily assume that \mathcal{S} is such that each player i gets to omit at least one favor for at least one other player, so so $\mathcal{B}_i(\mathcal{S}) < \mathcal{C}_i$ for all i . Recall that $M_{ij}(\mathcal{S})$ is the most costly favor that i must do for j in the selection \mathcal{S} ; recall as well that m_{ij} is the most costly favor that i does for j if i is doing all favors.

Then, in the *mild-punishment specification*:

1. Players are at any point in the game in one of two states (or everyone is in autarky). The two states are as always G, for *grace*, and P, for *purgatory*. All players begin in state G.
2. Player i in state G does all i -for- j favors in the selection \mathcal{S} for all j . A player i in state P does *all* i -for- j favors, for all j .
3. A player in state G, if she fails to do a favor, is forgiven by the immediate favor receiver j if the favor is not in \mathcal{S}_{ij} , and play continues with no additional consequences.
4. But if a player in state G fails to do a favor from \mathcal{S}_{ij} , she is not forgiven by the corresponding j , and a public randomization is conducted with two possible outcomes: autarky for all, forever; or i is sent to state P.
5. If a player is in state P and fails to do any favor, everyone moves immediately to autarky forever. (That is, forgiveness is *not* possible for transgressions committed when in state P.)
6. When and if player i is placed in state P, a publicly observable, exponentially distributed random alarm clock is set, with rate ϕ_i (which may be zero); when this clock goes off, player i returns to state G.

In the *stronger-punishment specification*, pieces 1 through 6 are modified as follows:

- 1'. Players are either in state G or state P (or autarky). All players begin in state G.

- 2'. No one does any favors for a player j who is (currently) in state P, and the failure to do a favor for a player j currently in state P is treated as if the favor had been done. If player i is in state G, she does all i -for- j favors called for by \mathcal{S} , for all j who are in state G; if i is in state P, she does all i -for- j favors (called for by \mathcal{S} or otherwise), for all j who are in G.
- 3'. If player i is in state G and she fails to do a favor for some j who is in G, she is forgiven by the immediate favor receiver j if the favor is not in \mathcal{S}_{ij} , and play continues with no additional consequences.
- 4'. But if i in state G fails to do a favor for some j who is in G that is called for under \mathcal{S} , she is not forgiven by the corresponding j , and a public randomization is conducted with two possible outcomes: autarky for all, forever; or i is sent to state P.
- 5'. If player i is in state P and fails to do a favor for some j who is in G, everyone moves immediately to autarky forever.
- 6'. When and if player i is placed in state P, a publicly observable, exponentially distributed random alarm clock is set, with rate ϕ_i ; when this clock goes off, player i returns to state G.

Please note that in both specifications, we have left unspecified the randomizing probabilities used in steps 4 and 4', respectively, and we treat the “return from state P” rates ϕ_i as parameters.

Proposition OA3. *For the game in which the favor receiver can issue ex-post cheap talk, the mild-punishment strategy profiles given by 1 through 6 above constitute a perfect social equilibrium for an appropriate choice of randomizing probabilities in part 4 if, for each i ,*

$$\max_{j \neq i} m_{ij} \leq \frac{1}{r} \frac{r(\mathcal{A}_i(\mathcal{S}) - \mathcal{C}_i) + \phi_i(\mathcal{A}_i(\mathcal{S}) - \mathcal{B}_i(\mathcal{S}))}{r + \phi_i}, \quad (\text{MP-ICP})$$

$$\text{and } \max_{j \neq i} M_{ij}(\mathcal{S}) \leq \frac{\mathcal{C}_i - \mathcal{B}_i(\mathcal{S})}{r + \phi_i}. \quad (\text{MP-ICG})$$

And the stronger-punishment strategy profiles 1' through 6' constitute a perfect social equilibrium for the appropriate randomizing probabilities if, for each i ,

$$\max_{j \neq i} m_{ij} \leq \frac{1}{r} \frac{r(0 - \mathcal{C}_i) + \phi_i(\mathcal{A}_i(\mathcal{S}) - \mathcal{B}_i(\mathcal{S}))}{r + \phi_i}, \quad (\text{SP-ICP})$$

$$\text{and } \max_{j \neq i} M_{ij}(\mathcal{S}) \leq \frac{\mathcal{A}_i(\mathcal{S})}{r + \phi_i}. \quad (\text{SP-ICG})$$

(The labels on the four inequalities are explained below.)

Proof. It is relatively straightforward to see that the strategies described are social; the only time that private information about the cost and benefit of a particular favor opportunity is used in any player’s strategy is at the moment of that favor opportunity.

To verify that these are (perfect) equilibria, if the appropriate inequalities hold, we must check that no player will wish to deviate at any single point in time, from any initial starting position in the game. Actions are only taken at the moment that favor opportunities arise; and a “starting position” (given the memoryless property of all the exponentially distributed times in the formulation and the nature of the prescribed strategies) is specified by a list of which players are in state G and which are in state P. Also, because this is a discounted formulation with bounded reward flows, we need only check on one-time deviations by any player, where the player assumes that, after the momentary deviation, she will go back to conforming to the allegedly equilibrium strategy.

When an i -for- j favor opportunity is not done by an i who is in state G, the favor receiver j must be indifferent between forgiving the favor or not doing so. We know that we can do this with the appropriate choice of randomizing probabilities in steps 4 and 4’, once we show that sending i to state P is (weakly) better for j than not doing so, and not doing so is better than autarky. We will show that continuation expected payoffs in state G exceed those in state P, and those in state P exceed autarky, so it is clear that, for j , continuation payoffs if he doesn’t forgive are better than the autarkic payoff 0. As for the assertion that j is (weakly) better off by sending i to state P than allowing her to remain in G: In the mild-punishment regime, the impact on j of sending i to state P (versus allowing her to remain in state G by forgiving her) is that i must do weakly more favors, at least for a while, which is of course weakly better for j . And in the stronger-punishment regime, j benefits (weakly) on those grounds and, in addition, will for a while be able to omit all favors for i , also weakly better for j .

What remains, then, is to show that i , if in state G, would prefer to do all favors called for under \mathcal{S} , to avoid being sent to state P (or, worse, to autarky), and that if she is in state P, she is willing to do all favors, to avoid triggering autarky. In both the mild-punishment and the stronger-punishment regimes, the first of the two inequalities in the displays guarantees the latter condition, and the second inequality guarantees the former.

Reason as follows: Along the path of play, starting from any position, players are never sent to state P; they serve out their time in that state if that is their initial position, and they remain in G once they get there or if that is their initial position. A defection from the alleged equilibrium strategies triggers autarky if i is in P, so to show that i is willing to do all favors when in state P, we must show that i ’s continuation value in P beats the continuation value of autarky, which is 0, by at least $\max_{j \neq i} m_{ij}$.

In the mild-punishment regime, the expected flow of payoffs to i if she is in state P (and hereafter conforms to the equilibrium strategy) is at least $\mathcal{A}_i(\mathcal{S}) - C_i$ for as long as she is in state P, and it improves to at least $\mathcal{A}_i(\mathcal{S}) - B_i(\mathcal{S})$ once she escapes. (In the mild-punishment regime, i receives at least all the favors specified by \mathcal{S} —she might receive more if others are in state P for a while—and she must do all favors while in P and those specified by \mathcal{S} once she escapes.) So, conditioning and unconditioning on

the time t she escapes from P, a lower bound on her continuation payoff in state P is

$$\int_0^\infty \left[\int_0^t e^{-rs} (\mathcal{A}_i(\mathcal{S}) - \mathcal{C}_i) ds + \int_t^\infty e^{-rs} (\mathcal{A}_i(\mathcal{S}) - \mathcal{B}_i(\mathcal{S})) ds \right] \phi_i e^{-\phi_i t} dt,$$

which is easily shown to be $[r(\mathcal{A}_i(\mathcal{S}) - \mathcal{C}_i) + \phi_i(\mathcal{A}_i(\mathcal{S}) - \mathcal{B}_i(\mathcal{S}))]/[r(r + \phi_i)]$. Hence, inequality (MP-ICP) (for *Mild Punishment–Incentive Compatibility in state P*) ensures that i , if in P, will do all favors. And in the stronger-punishment regime, we employ the same logic, although the expected flow of benefits to i while in state P is 0 (since no one is doing any favors for her), hence in place of the $(\mathcal{A}_i(\mathcal{S}) - \mathcal{C}_i)$ in (MP-ICP), we have $(0 - \mathcal{C}_i)$ in (SP-ICP) (where SP stands for *stronger punishment*).

And, to ensure that player i , if in state G, will do all favors called for under \mathcal{S} , it suffices to show that her continuation value in state G exceeds her continuation value in state P by at least $\max_{j \neq i} M_{ij}(\mathcal{S})(\mathcal{S})$. In the mild-punishment regime, the difference in continuation values is the difference in the extra favors she must do, for as long as she is in P, which is

$$\int_0^\infty \left[\int_0^t e^{-rs} (\mathcal{C}_i - \mathcal{B}_i(\mathcal{S})) ds \right] \phi_i e^{-\phi_i t} dt = \frac{\mathcal{C}_i - \mathcal{B}_i(\mathcal{S})}{r + \phi_i}.$$

For the stronger-punishment regime, she must do (weakly) more favors in P than in G, however how many more (and for how long) depends on who else begins in P. It is clear, though, that she loses the value of favors done for her, for as long as she remains in P, so an underestimate of the difference in the two continuation values is the expected loss of the benefits of those favors, which by a similar calculation is $\mathcal{A}_i(\mathcal{S})/(r + \phi_i)$. Therefore, the inequalities (MP-ICG) and (SP-ICG) (where ICG stands for *Incentive Compatibility in state G*) guarantee that i in state G will do all favors called for under \mathcal{S} , in the two respective regimes. ■

That completes the proof that the two regimes constitute perfect social equilibria (if the corresponding inequalities hold). But it is worth making a few more remarks.

Note first that in the stronger-punishment regime, the first of the two inequalities cannot possibly hold unless $\phi_i > 0$, while in the mild-punishment regime, $\phi_i = 0$ might work. This is the tip of an important intuitive iceberg: For this equilibrium to work, player j must weakly prefer to send player i to state P; this is needed to ensure that, with the appropriate mixture with autarky, j is indifferent between forgiving i and not. But, then, state P must be a state in which i does favors for others or, at least, for j . This in turn means that, in state P, i must have incentive to do favors. In the mild-punishment regime, i continues to receive favors; if the flow rate of favors received (which is at least) $\mathcal{A}_i(\mathcal{S})$ exceeds the flow rate of all favors given \mathcal{C}_i , i is getting a positive flow of overall value; if r is close enough to zero, that provides i with the incentive needed, even if i can never escape P (that is, if $\phi_i = 0$). But in the stronger-punishment regime, i receives no favors in state P, and only the prospect of an escape from state P (which means $\phi_i > 0$) can give i the incentive to perform while in i .

What happens if $\mathcal{B}_i(\mathcal{S}) = \mathcal{C}_i$ for some i ? For such an i , \mathcal{S} calls for her to do all favors, and some special handling is required. Without going into details, for any such player, one can invent a third state, call it P^* , to which all such i are assigned *at the start and permanently thereafter*. A player in this state (we might say instead, of this sort) is required to do all favors; if she fails, she triggers autarky for all. But even in the stronger-punishment regime, there is no reason not to do favors for such players; in terms of favors *for* players in P^* , treat them as if they were in state G .

This special handling of players for whom $\mathcal{B}_i(\mathcal{S}) = \mathcal{C}_i$ is indicative of a further point: We have described two possible perfect social equilibrium regimes, and the paper gives a third, but there are many other candidates. In particular, these two can be “mixed and matched” on a player-by-player basis, or even on a player-pair-by-player-pair basis. That is, we can imagine possible equilibria in which some players are subjected to mild-form punishment while others come in for stronger-punishment, or even where the form of punishment is based on the identity of the player who is not forgiven as well as on the identity of the player who did not forgive. What works (best or at all) will depend on the exact parameters of the situation.

Put it this way: The typical construction in “long horizon yields cooperation” games involves an equilibrium that is valuable enough to each player, so that players will keep to the cooperative scheme for fear of losing that value. In this construction, we require players to confirm reliably and publicly that others are conforming to the cooperative scheme, based on things they know but that are not common knowledge. To ensure reliability in these cheap-talk reports requires that the reporting player is unaffected by what he says. Therefore, if the cheap-talk report is going to cause someone else to be “punished,” that punishment cannot materially harm the reporting player. Therefore, we construct “purgatory,” in which a malfeasant can make up for her malfeasance, in a way that weakly improves the situation for others; she must (at least) do all the favors that she would have done had she not malfeasted. In the equilibrium of the paper, she is given the incentive to do those (and only those) favors by the promise of a return to grace, if she does them. (But, to make things work, the return to grace is with a probability that renders her indifferent between doing these favor or not.) Here, she is required to do all favors, under the threat of autarky, and so the equilibrium value in these constructions must be divided over two incentive constraints, a constraint that ensures a player in state P will perform if the alternative is something worse; and a constraint that state P is bad enough relative to state G , so the player will act as desired in state G .

The two equilibrium regimes, by giving us a variety of forms for state P , give us a better chance of structuring P so that both incentive constraints hold. Mild punishment makes state P not too unattractive relative to G (and the chance to escape makes P even less unattractive), giving more room between P and autarky. Stronger punishment allows us to make P more unattractive if that is needed, although (of course) doing so makes it more difficult to maintain the required incentive constraints when in P .

Note finally that these equilibria deliver on the promise made at the start: If, for a particular selection \mathcal{S} and interest rate r , the inequalities in the proposition are strict, then every choice made by the favor giver is a strict best choice; and the only time the favor receiver is required to speak is when a favor giver

who is in state G fails to do a favor; favor receivers play no role in restoring others to state G.

It goes without saying that there will be cases for which the equilibrium of Proposition 2 in the paper works, and the inequalities that are needed for these two constructions fail. But for a given selection \mathcal{S} , they either all work or none do, for “small enough” interest rates.

Corollary to Proposition OA3. *Suppose that $\mathcal{A}_i(\mathcal{S}) > \mathcal{B}_i(\mathcal{S})$ for all i that are called upon to do a favor under \mathcal{S} . Then for all small enough r , both strategy profiles in Proposition OA3 (for appropriate ϕ_i), as well as the strategy profile in Proposition 2 of the paper, constitute perfect social equilibria. Conversely, suppose that $\mathcal{A}_i(\mathcal{S}) \leq \mathcal{B}_i(\mathcal{S})$ for some i who is called upon to do a favor under \mathcal{S} . Then \mathcal{S} cannot be implemented in any equilibrium (whether perfect or social or not), with any sort of cheap talk.²*

The corollary doesn’t explicitly assume that $\mathcal{C}_i > \mathcal{B}_i(\mathcal{S})$ for all i . If $\mathcal{C}_i = \mathcal{B}_i(\mathcal{S})$ for some i , those i are handled with a third state P* as discussed previously. Let $\phi_i = \sqrt{r}$ for all other i . (Any relation between the ϕ_i and r will work as long as $r/\phi_i \rightarrow 0$ and $\phi_i \rightarrow 0$, both as $r \rightarrow 0$.) For those i such that $\mathcal{C}_i > \mathcal{B}_i(\mathcal{S})$, the right-hand side of (WP-ICG) goes to infinity as $r \rightarrow 0$, so for all small r , this inequality holds. And since $\mathcal{A}_i(\mathcal{S}) > \mathcal{B}_i(\mathcal{S}) \geq 0$ for all i , the right-hand side of (SP-ICG) goes to infinity as $r \rightarrow 0$; this inequality holds for small r . And on the right-hand side of the two ICP inequalities, as $r \rightarrow 0$, both $[r(\mathcal{A}_i(\mathcal{S}) - \mathcal{C}_i) + \phi_i(\mathcal{A}_i(\mathcal{S}) - \mathcal{B}_i(\mathcal{S}))]/[r + \phi_i]$ and $[r(-\mathcal{C}_i) + \phi_i(\mathcal{A}_i(\mathcal{S}) - \mathcal{B}_i(\mathcal{S}))]/[r + \phi_i]$ approach $\mathcal{A}_i(\mathcal{S}) - \mathcal{B}_i(\mathcal{S}) > 0$, so when we divide as well by r , the entire right-hand side goes to infinity. Once more, for sufficiently small r , the inequality holds. And it is apparent that condition A of Proposition 2 in the paper holds for small enough r .

As for the second part, because we have assumed that the cost of every favor is strictly positive, if $\mathcal{A}_i(\mathcal{S}) \leq \mathcal{B}_i(\mathcal{S})$ for some i who must do a favor according to \mathcal{S} , she will not do this favor when called upon to do so in any equilibrium that allegedly implements \mathcal{S} . ■

Variations on the information endowments

Throughout this analysis, we have assumed that, for each i -for- j favor with cost–benefit vector (x, y) , both i and j know both x and y perfectly. But what if one or the other or both know something less? We have in mind the following general structure: When an i -for- j favor arises, and dependent on the true values of cost and benefit, i and j separately receive (possibly different) information about the values of x and y . We always assume that i learns at least x and j learns at least y .³ This information, which we call their (respective) information endowments, arrives before any communication can take place between them and before i must decide whether to do the favor; moreover they receive no further

² To be very clear, the term *implemented* here means: Along the (equilibrium) path of play, the favors that are done are those in \mathcal{S} and only those in \mathcal{S} .

³ This is without loss of generality as long as we make the weaker assumption that i has (weakly) superior information (to j) about her costs, and j has (weakly) superior information about his benefits, and we interpret x as the conditional expectation of costs given i ’s information and y as j ’s conditional expectation of benefits.

information except as is provided via communication between them (which others may share).

In this setting, the equilibria of Proposition 1 and OA3 are robust in the following sense:

Proposition OA4. *Suppose that the information endowments are such that the favor giver in all cases assesses probability at least $1 - \epsilon$ that the favor receiver's benefit is some \hat{y} , and she assesses this for \hat{y} as the true y with probability at least $1 - \epsilon$. And, in the same circumstances, the favor receiver (who knows y) in all cases assesses probability at least $1 - \epsilon$ that the favor giver's cost level is some \hat{x} , and he assesses this for \hat{x} as the true x with probability at least $1 - \epsilon$. Fix a selection \mathcal{S} and interest rate r . Then, if, for the equilibrium of Proposition 1 (in the paper) or either equilibrium in Proposition A3, the required inequalities all hold strictly, there is some $\epsilon^* > 0$ (depending on r and the other data of the game) such that for $\epsilon < \epsilon^*$, the corresponding behavior strategies, adapted so that players act according to their \hat{y} or \hat{x} "predictions" and, for the two equilibria in Proposition A3, with the randomizing probability used to render the favor receiver indifferent suitably adapted, give a perfect social equilibrium with payoffs that are close to $(A_i(\mathcal{S}) - B_i(\mathcal{S}))/r$ for the different players i . And as $\epsilon \rightarrow 0$, the equilibrium payoff for i approaches $(A_i(\mathcal{S}) - B_i(\mathcal{S}))/r$.*

Before discussing the proof, let us indicate a situation in which the proposition's assumption on players' posterior beliefs on y and x will hold: Suppose that, when a favor opportunity arises, with cost-benefit vector (x, y) , the favor giver i is provided with a signal that is one value from the support of y . With probability $1 - \delta$, the signal is the true value of y , while with probability δ , it is some mixture of the other elements of the support of y . And the favor receiver j is provided with a signal drawn from the support of x that is the true value of x with probability $1 - \delta$ and some other value from the support of x with probability δ . Then, based solely on these signals, as $\delta \rightarrow 0$, Bayes' rule leads i and j to assess that the truth is very likely the signal they receive, with probability that approaches 1. Knowledge of their own part of the payoff can only improve this, although one has to be careful here: For some (x, y) pairs, knowledge of x may reveal the value of y . In such cases, even if i 's signal is some $y' \neq y$, i will assess probability 1 that that the benefit level is the true y , and this is the \hat{y} in the proposition.

The proof of the proposition is simple: As long as the required inequalities from the propositions hold strictly, all actions by the players (except for j telling the "truth") are strict best responses when players know x and y precisely. Here, there is a bit of noise, which shifts continuation values and continuation payoffs slightly, but with strict inequalities, small changes don't affect the strict optimality. The only actions that are finely balanced in the equilibria are j 's declaration of forgiveness (or not) in the two equilibria of Proposition OA3, if i doesn't do a favor from \mathcal{S} , and as we can make j indifferent by shifting the randomizing probability a bit as needed, we can keep j indifferent and, therefore, telling the "truth" as he perceives it (that is, if he assesses high probability of some \hat{x} , he forgives i if \hat{x} is less than the true y). Note that, for positive ϵ , in the two equilibria of Proposition OA3, players will wind up in state P and, in the equilibrium of Proposition 1 and perhaps in the two equilibria of Proposition OA3, eventually players will land in autarky with probability one. But the odds of this happening are

vanishingly small for any single incident, and so for fixed $r > 0$, they are vanishingly small in a time frame that has a nonnegligible impact on equilibrium payoffs.

What of local information that is far from perfect? Under the assumption that i knows x and j knows y , the furthest we can be from perfect local information is the case where the favor giver i knows only x and the favor receiver knows only y .⁴ Even in this case, the ability of i and j to communicate through cheap talk can enhance their payoffs. Consider the special case in which the cost x and benefit y of any favor (for any ordered pair (i, j)) are probabilistically independent:

Proposition OA5. *Suppose that, in all cases, a favor giver knows only her own cost x while the favor receiver knows only his own benefit y . Suppose that, in all cases, the distribution of benefits and costs (for each ordered pair i and j) is such that x and y are probabilistically independent. And suppose that, for each player i , selection \mathcal{S} generates an expected overall payoff greater than 0, or $A_i(\mathcal{S}) > B_i(\mathcal{S})$ for all i . Then as $r \rightarrow 0$, perfect social equilibria can be constructed (for a game that is augmented by cheap-talk announcements) whose normalized payoffs for each i approach $A_i(\mathcal{S}) - B_i(\mathcal{S})$.*

This result is more or less a direct corollary of the general results of Fudenberg, Levine, and Maskin (1994, Section 8).⁵ It is interesting to speculate on whether similar results obtain (for r going to zero) if the players have *better* local information or if the benefit–cost vectors are not independent, but we will not chase down results of this sort here.

Instead, we focus on results for a given level of $r > 0$. Proposition OA5, combined with Proposition OA2, tells us that even if i knows only x and j knows only y , the ability to communicate with cheap talk may improve matters for the players, for some $r > 0$: Per OA2, if the players cannot engage in cheap talk, then there are some parameterizations such that the u-efficient selection \mathcal{S}^U cannot be attained, even asymptotically. Hence the ability to engage in cheap talk about this (far from perfect) local information does improve matters, at least in some situations and for some $r > 0$.

At this stage, it would be interesting to launch into a full-fledged analysis of different forms of (less than perfect) local information. Some special cases are particularly interesting; for instance, what if the favor giver and receiver both know the cost of every favor (but only the receiver knows his benefit), or if they both know the benefit level (but only the favor giver knows her cost)? In the paper, we already picked off some low lying fruit along these lines. In particular, we mentioned in passing the following result, which we (now) formalize.

Proposition OA6. *For a given selection \mathcal{S} , suppose that when an i -for- j favor arrives, i and j know (perhaps) less than the values of the cost x and the benefit y , but they both know whether $(x, y) \in \mathcal{S}$ or not. Then the equilibria of Proposition 1 and of Proposition OA3 continue to be equilibria, when the required inequalities hold.*

⁴ Of course, knowledge of x may entail some knowledge of y and vice versa, depending on their supports and joint distribution; we'll discuss this below.

⁵ We say “more or less” because in our setting, the timing of actions is random. But this poses no serious problems.

The proof is standard: If a player knows that a particular action is her best response, she doesn't need finer information. This cannot be applied to the equilibria of Proposition 2, because in that equilibrium, there are times when j must know x (so that j can conduct the appropriate randomization).

In view of this, what can be said about cases in which either i or j do not know whether a given favor is in the selection \mathcal{S} ? We conclude with a proposition and then an example that show some of the complexities that can arise.

In doing so, we are fairly agnostic as to the rules of the interaction between i and j at the moment of an i -for- j favor. They may engage in cheap talk both before and/or after i must choose whether to do the favor or not, with the cheap talk held privately between them or broadcast to all players. But we maintain the following structural assumptions: (1) Rules for what is communicated must be set in advance and adhered to, although the communication can be contingent on earlier communication or on i 's decision whether to do the favor. (2) The structure of the informational endowments of the players are common knowledge: when an i -for- j favor opportunity occurs at some point, (no more than) the past public history of the game determines the structures of what i and j are told (but not, of course, their specific signals they receive), and these structures are common knowledge. (3) Player i must choose between doing the favor or not. And (4) all other interactions involve only cheap talk.

Proposition OA7. *Suppose that, for some ordered pair i and j , there is an i -for- j favor at which either i or j , based on her or his information endowment, assesses positive probability for both $(x, y) \in \mathcal{S}$ and $(x, y) \notin \mathcal{S}$. And suppose that, in this event, the other player has superior information to the first. Then, regardless of any arrangement of cheap talk, the selection \mathcal{S} cannot be implemented in a perfect social equilibrium for any $r > 0$.⁶*

Proof. Suppose \mathcal{S} could be implemented. Take the pair i and j for which the hypothesis is true, and suppose it is i that assesses positive probability for both $(x, y) \in \mathcal{S}$ and $(x, y) \notin \mathcal{S}$, at which event j has superior information. Because j has superior information, he knows that i assigns positive probability to these events. To implement \mathcal{S} , it must be that the favor is done if $(x, y) \in \mathcal{S}$ and not done if $(x, y) \notin \mathcal{S}$, and the continuation from these separate immediate actions is the continuation of implementing \mathcal{S} . Since j 's information is superior to i 's, for this to happen, j must know (be endowed with) whether $(x, y) \in \mathcal{S}$ or not. And if j knows that $(x, y) \notin \mathcal{S}$, he can act as if he knew instead that $(x, y) \in \mathcal{S}$, have the favor done for him, and continue with the implementation of \mathcal{S} , which is better for him than if the favor is not done. That is, j is in a position to manipulate i into doing the favor "for free" (in terms of any subsequent consequences).

On the other hand, if it is j who, in the circumstance in question, assesses positive posterior probability to both $(x, y) \in \mathcal{S}$ and $(x, y) \notin \mathcal{S}$, and i who has the superior information, then the argument just given can be refashioned to imply that i knows which of the two situations prevails and knows that

⁶ This does *not* say that we cannot approach the payoffs of selection \mathcal{S} asymptotically, as $r \rightarrow 0$. This concerns implementing \mathcal{S} exactly, for fixed $r > 0$.

j doesn't know this; when $(x, y) \in \mathcal{S}$, i can act as if she instead knew that $(x, y) \notin \mathcal{S}$, causing the favor not to be done with no untoward consequences in terms of the continuation; she can manipulate him into letting her escape doing the favor "for free."

In either case, the selection \mathcal{S} is not implemented, which is a contradiction. ■

The first antecedent assumption in Proposition OA7, that some circumstance arises in which one of the players is (conditionally) uncertain whether $(x, y) \in \mathcal{S}$ or $(x, y) \notin \mathcal{S}$, is the complementary condition to the antecedent assumption in Proposition OA6. To that degree, the two are converse. But, of course, the second part of the antecedent assumption in Proposition OA7—that in some circumstance where a player, based on his or her endowment, is unsure whether the favor should be done or not, the other player knows this—is rather strong. It does hold in some "natural" conditions; e.g., if j knows x as well as y , or if i knows y as well as x . But something like this assumption is needed: The idea in the proof is that when a player can't tell whether the favor is to be done or not, the other player can manipulate the first in either direction and, of course, chooses to manipulate in the direction the second player desires. Without this sure knowledge that the other player can be manipulated, we might be able to implement \mathcal{S} . A simple two-player example, constructed for the u-efficient selection \mathcal{S}^U , shows this:

Call the two players Alice and Bob. Alice-for-Bob and Bob-for-Alice favors each arrive at the rate 1, and the distributions of costs and benefits are the same: The favor giver always has a cost of 5, while the favor recipient's benefit is 4 with probability 0.1 and 10 with probability 0.9. If Alice is giving the favor, both Alice and Bob know both the benefit level and cost. But if Bob is giving the favor (and Alice is receiving it), Bob's information endowment is more complex: he learns her benefit level with probability 0.9, but doesn't learn it with probability 0.1 (independent of all other events). And Alice is unsure whether Bob has learned her benefit level or not.

The following implements the u-efficient outcome for small enough r . Regarding Alice-for-Bob favors, we follow the behavior rules set forth in the mild-punishment regime of Proposition OA3.⁷ And, regarding Bob-for-Alice favors, Alice must declare her benefit level before Bob does the favor (or not). Her equilibrium strategy is to declare the truth. If she declares a benefit of 10 and Bob either knows that this is correct or doesn't know, he does the favor. If she declares a benefit of 4, Bob doesn't do the favor. The key is what happens if she declares 10 when Bob knows the benefit level is 4: Bob doesn't do the favor and instead declares that he knows she lied, and a public randomization takes place at which either autarky immediately ensues or she is sent to a state where she must do all favors for Bob, forever, on peril of causing autarky. The randomizing probability is chosen so Bob is indifferent between declaring her a liar or doing the favor; note that he prefers her in this punishment state to not having her there to autarky, so a randomizing probability can be found that renders him indifferent. We leave it to the reader to verify that for small r , this all holds together.

⁷ Since there are only two players, and Alice and Bob both know x and y for these favors, we can simplify: Bob needn't forgive Alice if she doesn't do a favor of benefit level 4; and if she doesn't do a favor of benefit level 10, the players can immediately move to autarky. But it doesn't hurt to follow the behavior rules from Proposition OA3 exactly.

The example “works” because in the cases where Bob doesn’t know whether to do the favor, Alice is unsure whether Bob is uncertain. Hence, she cannot safely manipulate him into thinking that he should do the favor, when (she knows) he should not. The second antecedent assumption in Proposition OA7 is, of course, violated.

On the other hand, while the first antecedent assumption of Proposition OA7 must hold, the second is not necessary: Suppose that for a particular ordered pair i and j , the distribution of benefit–cost vectors (x, y) is such that knowledge of x does not preclude any y in the support and vice versa. (Suppose, for instance, that they are probabilistically independent). Suppose that i knows only x and j knows only y . Then the u-efficient outcome cannot be implemented. Suppose by way of contradiction that it could. Then if the highest value of x is larger than the highest value of y , j would have to be able to force i to do the favor whenever he sees that y is at its highest value. But then he could mimic his “highest- y ” behavior for any other y and have the favor be done. And if the highest value of x is higher than the highest value of y , then i is able not to do the favor when she sees the highest value of x with no untoward consequences. But then by mimicking her “highest- x ” behavior for any other x , she can avoid doing the favor.