Rules With Discretion and Local Information*

T. Renee Bowen
David Kreps
Andy Skrzypacz

August 2012

Abstract

To ensure that individual actors take certain actions, community enforcement may be required. This can present a rules-versus-discretion dilemma: It can become impossible to employ discretion based on information that is not widely held, because the wider community is unable to tell whether the information was used correctly. Instead, actions may need to conform to simple and widely verifiable rules. We study when discretion in the form of permitted exceptions to the simple rule can be permitted, if the information is shared by the action taker and a second party, who is able to verify for the larger group that an exception is warranted. In particular, we compare protocols where the second party *excuses* the action taker from taking the action *ex ante* with protocols where the second party instead *forgives* a rule-breaking actor *ex post*, finding that the latter is, in general, useful in a wider variety of circumstances.

*Comments by seminar participants at UC Santa Barbara and Stanford Graduate School of Business are gratefully acknowledged, as is the financial support of the Stanford Graduate School of Business.

1 Introduction

Economists have long advanced the notion that, in some contexts, decision makers should forbear from using their best judgement and instead follow some rule. The context of central banking is perhaps the first in which this idea was explored, by Simons (1936) and subsequently (for instance) by Modigliani (1977) and Lucas (1980). Other prominent contexts include financial accounting (Bratton [2003]; Schipper [2003]; Barth [2008]), the practice of medicine (Kessler [2011]), and jurisprudence (Freed [1992]; Becker [1997]).

The argument for the use of discretion typically comes down to (superior) information held by the decision maker; the counter-arguments in favor of the application of a rule vary, but among them (and relevant to this paper) are the need for *ex post* verifiability by a wider group of individuals that the decision taken was "appropriate." In particular, where one depends on *community enforcement* of decisions made by individuals, it may be necessary that individuals in the community can perform this *ex post* verification.

While much of the literature has stressed a dichotomous choice between rules and discretion, it will come as no surprise to an observer of the real world that the matter is more nuanced: Decision makers are given limited amounts of discretion in some cases; in other cases rules are established, but exceptions are permitted and/or violations to the rule are forgiven.

For instance, Bowen (2011) examines the decision-making and enforcement procedures of the General Agreement on Tariffs and Trade (GATT). (For related work, see also Maggi [1999].) Under GATT, general rules are applied on a global basis to the trading relations between pairs of countries that, broadly speaking, say that departures from free trade should be punished (with a retaliatory tariff). However, as Bowen observes, numerous departures from free trade go unpunished. That is, country A may exercise some discretion in applying a retaliatory tariff to country B when country A sympathizes with the reason country B departed from free trade. Such an equilibrium may not be possible with only bilateral relationships, because it is very costly to withhold tariff retaliation. But the threat of the GATT agreement falling apart (multilaterally) provides enough incentive to do this.

Or consider the case of the Toyota system of subcontracting, in which Toyota forms long-term "strategic partnerships" with firms that supply Toyota with major sub-assemblies (see Milgrom and Roberts [1992]; Kreps [2004, Chapter 24]). Toyota, acting as a hierarchical superior, more or less dictates terms (prices to be paid, quantities to be ordered, designs to be followed) to its strategic partners, who moreover are expected to (and do) make sunk-cost investments in the relationship. Having made these sunk-cost investments, the suppliers are protected from hold-up exploitation by Toyota largely through the threat of the collective action of the all of Toyota's suppliers, to be triggered if any one of them is unfairly treated by Toyota. Faced with the collective power of its suppliers, Toyota has a powerful incentive to maintain its reputation for fair-dealing with individual suppliers. Toyota follows general rules, so that suppliers can verify that Toyota has dealt fairly in individual cases. But exceptions to these rules are sometimes made, and Toyota expects the affected supplier to reassure its peers that the exception was legitimate; Toyota even provides the forum—various industry groups of its suppliers—at which such reassurances can be offered.

Or think of the case of the faculty of a school or department. Rules are adopted for how individual faculty members behave, what are their responsibilities, and so forth. But deans and department chairs, acting on behalf of the school or department, will sometimes allow exceptions to or forgive transgressions against these rules in specific cases, expecting that members of the department/school will trust that the exceptions/forgiveness are allowed for appropriate reasons.

We contend that in all these and in similar cases, a key is that the "private information" on which basis the exception is made is not private in the strict sense that only one party (the party who does not follow the rule) has access to the information, but instead is *local*, by which we mean: held by *both* the active party who does not follow the rule and by the second party whose (lack of) action reassures members of the larger community. [1] So, for instance, we expect that exceptions to rules are likely to be more prevalent in circumstances where the second party is more likely to be able to judge if an exception was appropriate; i.e., where information is local rather than private. Casual empiricism suggests that this prediction is true; certainly it seems the case that, in the context of GATT enforcement, departures from free trade are more likely to go unpunished the closer the two countries are diplomatically, which presumably means the more likely it is that the offended party will understand why the offending party did what it did.[2]

Indeed, we contend that, in these sorts of situations, parties expend resources to

[1]Our use of the adjective "local" echoes the way the term is used in Wolitzky (2012). But Wolitzky is concerned with local information about actions the players take, while our concern is with local information about payoff-relevant but exogenous factors.

[2]Of course, diplomatic "closeness" suggests other explanations for why the offended country doesn't engage in retaliation. For one thing, one can imagine that the leadership of the offended country simply wishes to be more lenient with allies than with countries with which it has a more adversarial relationship. And to the extent that "closeness" is correlated with a greater level of trade relationships, the countries may be engaging in bilateral reciprocity; see the theoretical developments following.

make private information local: In the case of Toyota, for instance, exceptions are made on the basis of production conditions facing both the subcontractor and Toyota, and (as is well documented) Toyota expends significant resources—and expects its subcontractors to expend significant resources—so that each side of the transaction knows a great deal about the production conditions affecting the other party. Part of the duties of a department chair or dean—and we believe an important success factor for a chair or dean—is a willingness to understand the personal conditions that affect faculty members in the department or school. (And, of course, the same observation can be made about supervisors of work groups in general.) Now, of course, in both these cases, the information investment can be attributed in the first place as being driven by the desire to make "appropriate" decisions. But, we contend, a second reason is so that the second party can be a credible witness about whether the violation of some rule is legitimate.

To the best of our knowledge, the use of local information in this sort of situation has not been formally modeled. (But see the literature review following.) So the principal purpose of this paper is to provide a formal, stylized model of how local information might be employed. We verify the basic premise—that when information is local, a second party who has the information can be an effective and credible source in excusing or forgiving exceptions to rules based on (otherwise) private information—but go on from showing that this might be done to a more detailed examination of how. The paper runs as follows:

1. The basic model is set up and analyzed in Sections 2 through 4. Leaving details for later, the basic structure is an assembly of players, where players sometimes have the opportunity to do a favor for another player. The players directly involved (the favor giver and favor receiver) know the levels of cost and benefit of the favor; the other players observe only that the opportunity for a favor takes place and whether the favor giver chooses to do the favor. Even if the pattern and frequency of favor opportunities is such that pairs of players cannot "trade" favors, community enforcement may be effective. High-cost, low-benefit favors may be inefficient, but as the community cannot directly tell which favors should not be done, if community enforcement is needed, the "cost" of having favors done is that inefficient favors may have to be done. This is our rules versus discretion conflict.

2. In Section 5, we briefly consider unilateral discretion exercised by favor givers. This may enhance efficiency but can never lead to a completely efficient equilibrium; under some conditions, it can produce asymptotically (in the interest rate) efficient outcomes, but in many cases of interest, equilibrium payoffs can

be uniformly bounded away from efficient outcomes.

3. Through Section 5, the assumption that the favor receiver knows the cost and benefit of each favor is not used; while both the favor giver and receiver know these values, only the favor giver (she) acts. In Section 6, we consider allowing the favor receiver (he) to issue (cheap-talk) *ex ante* excuses: Having seen the cost and benefit levels, he can tell the community that he doesn't expect this favor to be done. We show how this can be used to enhance efficiency (relative to the unilateral exercise of discretion); but in some cases, it is still impossible to achieve full efficiency, even asymptotically. Interestingly, the nature of how the favor receiver is treated in equilibrium shifts from Section 5: In Section 5, it is natural to compensate a potential favor receiver who does not receive a particular favor; here we must attend to the receiver's incentives to excuse favors, which means that we should (weakly) punish a favor receiver when he doesn't receive a favor.

4. Section 7 briefly considers allowing the favor giver and receiver to make simultaneous announcements, following the sort of thing done in Krishna and Morgan (2001). We also (briefly) discuss why we find this theoretical construct to be unconvincing.

5. In Section 8, we imagine that instead of excusing the favor giver from doing the favor *ex ante*, the favor receiver might publicly *forgive* the favor giver for favors that are not done. This sort of construction works in the strong sense that efficiency can be achieved exactly for small enough r (if the basic parameters of the problem meet some simple conditions), even in cases where *ex ante* excuses fail to achieve efficiency asymptotically. This is because, once the favor has or has not been done, its particular benefits (and costs) are irrelevant to the continuation of the game.

6. Sections 9 and 10 concern extensions and elaborations of the basic model. Through Section 8, we assume that each of the favor giver and receiver know the cost and benefit levels of the favor perfectly. Section 9 analyzes situations where the two together know this, but individually they may have less than perfect "local information." And Section 10 considers elaborations where the favor generates positive externalities for others in the community, where the favor giver and receiver can engage in under-the-table but inefficient transfers, and then where both of these are present. Neither externalities nor under-the-table transfers alone cause problems for our earlier constructions; but when

both are present, much of what we said earlier will not work. We conclude in Section 11.

It should be noted at the outset that the notions of "rule" and "exception" are fluid; one can take the semantic position that we are merely investigating the replacement of a simple rule with a more nuanced rule that involves the exercise of some discretion in the application of the simple rule. We grant this semantic point, but we still believe the equilibria and conclusions we draw from them are economically interesting.

Related Literature

We are unaware of papers that ask the same questions we are asking, although there are two strands of literature that are related and employ similar langugage.

A characterization of what we do here is that we are looking at the transmission of information from a small set of (two) players to a larger community. This in some ways echoes the substantial literature on random matching and the folk theorem. Much of this literature concerns the issues that arise when players who are matched do not automatically know what has been the history of play of the person against whom they are matched; in such circumstances, how can what happened in the past between pairs of players be adequately communicated to others, adequate (at least) to sustain folk-theorem like outcomes? This includes very early work by Rosenthal (1979) and, more representative of this literature, Kandori (1992) and Ellison (1994). The working paper by Wolitzky (2012) is of this type; his language is very close to ours. But despite some superficial similarities, there are substantial differences: this literature is concerned with issues of "hidden action" or moral hazard, where the past play of a player is not known by everyone. In our model, we assume that each player knows the (public-information) history of play by all the other individuals in the society; the issues that we address concern the ability to transmit (in credible form) information about exogenous variables that only some players hold.

The transmission of private information about payoff relevant variables via cheap talk is nearly as "old," beginning with the seminal paper of Crawford and Sobel (1982). Sobel (2011) gives a summary of this literature. Especially relevant to our stylized model is the literature on "trading favors" (Mobius [2001]; Hopenhayn and Hauser [2004]) that deals with issues of "hidden information" and, similar to the model we develop, the question of parties doing favors for one another. But most of this literature concerns one informed player transmitting information to a second player. Perhaps the closest to what we do here is work, beginning with Krishna and Morgan (2004), on cheap-talk communication by multiple informed parties; see Sobel (2011, Section 5) for a more complete bibliography.

2 A basic stylized model

We work with variations on the following stylized model.

There are I players, indexed by $i = 1, \ldots, I$.

Time is continuous, indexed by $t \in [0, \infty)$.[3] Opportunities arise at random for players to do favors for one another, generated by independent Poisson processes. The rate of arrival of opportunities for i to do a favor for j, hereafter called an *i-for-j favor*, is denoted by λ_{ij}. Let T_{ij}^n denote the arrival time of the nth opportunity for i to do a favor for j. (Under the assumptions of independent Poisson arrivals, there is zero probability that at any time t more than one favor opportunity takes place.)

Whenever an i-for-j favor opportunity occurs, the cost to i of doing this favor and its benefit to j are determined randomly, independently of the (inter-)arrival time at which it applies. Think of there being, for each ordered pair i and j, an i.i.d. sequence of pairs of real numbers $\{(x_{ij}^n, y_{ij}^n);\ n = 1, 2, \ldots\}$, each of these sequences fully independent of all other such sequences and independent of all the arrival times T_{ij}^n, such that at the time T_{ij}^n of the nth i-for-j favor opportunity, the cost of the favor to i is x_{ij}^n and the benefit j receives if this favor is done is y_{ij}^n. We are not assuming that the values of x_{ij}^n and y_{ij}^n are independent of one another; we allow the distribution of each cost–benefit vector (x_{ij}, y_{ij}) to depend on the ordered pair (i, j). We assume these distributions are such that:

- Each x_{ij} and y_{ij} is strictly positive with probability 1.
- The probability that $x_{ij} = y_{ij}$ is zero, for each i and j.
- Each x_{ij} and y_{ij} has finite support.

The last of these assumptions is made to simplify some proofs; but an assumption that the supports of x_{ij} and y_{ij} are bounded is essential to several of our key results. Of course, the last assumption ensures that costs and benefits have finite expectation; so let

$$a_{ij} := \mathbf{E}[y_{ij}] \quad \text{and} \quad b_{ij} := \mathbf{E}[x_{ij}].$$

Also, let m_{ij} be the highest value that x_{ij} takes on with positive probability.

Available actions (for the time being) are simple: If, at time t, an i-for-j favor opportunity occurs, then i must decide whether to do the favor or not.

[3] Anticipating a bit, we use continuous time with Poisson arrivals of opportunities to do favors, to avoid a situation in which one player is called upon to do more than one favor at any single point in time.

The key to our model is the distribution of information: We assume that every player knows when any i has the opportunity to do a favor for some j and, subsequently, whether i does that favor or not. But, if i has the opportunity to do a favor for j at time t, the cost-benefit values (x_{ij}, y_{ij}) for this favor are common knowledge between i and j and unknown to all the other players.

The payoff for each player is the infinite-horizon discounted sum of the value of favors received less the cost of favors given, with an instantaneous interest rate of r. That is, a cost or benefit incurred at time t is discounted by e^{-rt}. Each player seeks to maximize the expectation of her (infinite-horizon) payoff.

The u-efficient outcome

Within this setting, there are a variety of outcomes that are (Pareto) efficient. We focus attention on a specific efficient outcome, the outcome that maximizes the sum of payoffs to all players.[4]

Definition 1. *The* ***u-efficient*** *outcome is where the favors that are done are precisely those for which* $x_{ij} \leq y_{ij}$.

The "u" in u-efficiency is shorthand for "utilitarian." Note that since we assume there is zero probability that $x_{ij} = y_{ij}$, this outcome is essentially (in the probabilistic sense) unique. Given our interest in this particular outcome, we define

$$A_{ij} := \mathbf{E}[y_{ij} 1_{\{y_{ij} \geq x_{ij}\}}] \quad \text{and} \quad B_{ij} := \mathbf{E}[x_{ij} 1_{\{y_{ij} \geq x_{ij}\}}],$$

where $1_{\{\cdot\}}$ is the indicator function of the event in the subscript. We let M_{ij} denote the upper bound of the support of $x_{ij} 1_{\{y_{ij} \geq x_{ij}\}}$, the largest possible cost of a favor, if i only does favors for j whose cost (to i) is less than their value (to j). And we define

$$\mathcal{A}_i := \sum_{j \neq i} \lambda_{ji} A_{ji} \ , \quad \mathcal{B}_i := \sum_{j \neq i} \lambda_{ij} B_{ij} \ , \text{and} \quad \mathcal{C}_i := \sum_{j \neq i} \lambda_{ij} b_{ij}.$$

That is, $\mathcal{A}_i$ is the expected flow rate of benefits accruing to i, if all the u-efficient favors (and only the u-efficient favors) are done for her, $\mathcal{B}_i$ is the expected flow rate of costs she accrues, if she does all the u-efficient favors she "owes" to others, and $\mathcal{C}_i$ is the expected flow rate of costs she accrues if she does *favors* for other players that come her way. (This notation is much used in the remainder of the paper.) Since $b_{ij} \geq B_{ij}$ for all i and j, $\mathcal{C}_i \geq \mathcal{B}_i$ for all i, with equality only if every favor that i might be called upon to do is u-efficient.

[4] Focusing on this one efficient outcome simplifies the exposition and statement of our results, but it may seem limiting. On this point, see the final comment in Section 10.

3 Autarky and Bilateral Reciprocity

A variety of perfect equilibria can be constructed for the basic model.[5] Or, put more correctly, we can describe a variety of strategy profiles and give conditions for them to be perfect equilibria.

Autarky

The simplest perfect equilibrium, which holds in all situations, is *autarky*: no one does any favors for anyone else; and all players get a payoff of 0.

Bilateral reciprocity

Focus on a particular pair i and j. Imagine that this pair adopts the following rule concerning their dealings with one another. Regardless of anything else that happens (with third parties), i does all favors she can for j, and j reciprocates by doing all the favors he can for her, as long as each behaves in this fashion. If either fails to do a favor for the other, they never subsequently do favors for each other. This, of course, is just bilateral reciprocity between i and j as in the repeated prisoners' dilemma, enforced by "grim punishment," adapted to this setting. The expected flow of benefits to i of this arrangement is $\lambda_{ji}a_{ji}$ while the expected flow of costs is $\lambda_{ij}b_{ij}$, so this is equilibrium behavior for i as long as the immediate cost of doing any favor for j, which has upper bound m_{ij}, is less than the expected, discounted value of ongoing benefits less costs. That is, this constitutes (perfect) equilibrium behavior for i and j (in their bilateral relationship) as long as

$$m_{ij} \leq \frac{\lambda_{ji}a_{ji} - \lambda_{ij}b_{ij}}{r} \quad \text{and} \quad m_{ji} \leq \frac{\lambda_{ij}a_{ij} - \lambda_{ji}b_{ji}}{r}. \tag{1}$$

Of course, this is but one possible bilateral arrangement between i and j. For instance, if

$$M_{ij} \leq \frac{\lambda_{ji}A_{ji} - \lambda_{ij}B_{ij}}{r} \quad \text{and} \quad M_{ji} \leq \frac{\lambda_{ij}A_{ij} - \lambda_{ji}B_{ji}}{r}, \tag{2}$$

then i and j can sustain their (bilateral) part of the u-efficient outcome: Each does favors for the other whose cost is less than the benefit, forbears from doing other favors, and if ever one fails to do a favor whose cost is less than its benefit, both do no favors for the other, forever more.

And from these bilateral equilibrium arrangements, full equilibria can be constructed. Suppose that (2) holds for some pairs i and j but not others. Then a full

[5] By "perfect," we mean "perfect Bayes," although we won't be precise about this. In fact, we are fairly confident that the equilibria we define are all sequential—we'll say a bit more about this later—but we won't verify that the consistency criterion holds.

equilibrium is where i and j behave in the fashion just described with each other for those i and j such that (2) holds, while if (2) does not hold for i and j, they do no favors for one another. Since we have asssumed that the bilateral relationships are additively separable (in the sense that what i gets on net from her relationship with j is added to what she gets from her relationships with others and so is unaffected by those other relationships, unless how j treats her and how she treats j is changed by what else happens), any sort of mix-and-match of bilateral equilibria gives an overall equilibrium for the game.

4 Social Enforcement

We are interested not in these pasted-together bilateral equilibria, but instead in equilibria for situations in which bilateral arrangements are ineffective. Two examples illustrate the sort of situation we have in mind.

Frequent favors, but not with any given partner
Suppose that $I = 11$, $\lambda_{ij} = 0.1$ for all i and j, $r = 0.1$, and the joint distribution of each (x_{ij}, y_{ij}) pair is degenerate at the point $(2, 3)$. For any pair i and j, the left-hand side of (1) (or [2], they are the same) is 2, while the right-hand side is $(0.1 \cdot 3 - 0.1 \cdot 2)/0.1 = 1$. Bilateral arrangements between pairs i and j will not work, because they interact too infrequently relative to the interest rate r.

No one serves the person who serves her
For any $I \geq 3$, suppose that $\lambda_{ij} = 1$ if $j = i + 1$ and $\lambda_{ij} = 0$ if $j \neq i + 1$, where we interpret $I + 1$ as 1. In words, if we arrange the individuals in a circle, numbered $1, 2, \ldots, I$ as we go clockwise around the circle, each i gets the opportunity to do favors for (only) her clockwise neighbor and receives favors (potentially) from (only) her anti-clockwise neighbor. To finish the example, suppose $r = 0.1$ and each (x_{ij}, y_{ij}) pair is degenerate at (2,3), as in the previous example. The point is that, clearly, for no pair is bilateral reciprocity going to work; for pairs i and j where $j = i+1$ or vice versa, one of the two inequalities in (1) holds, but the other fails; and for all other pairs i and j, $\lambda_{ij} = \lambda_{ji} = 0$ (the two never interact), so in that sense, the inequalities are irrelevant.

While bilateral reciprocity fails, equilibria in which all parties do all favors can be constructed in both these examples, where i is impelled to do favors for others (as the opportunity to do so arises) based on a threat of *social punishment* or *community enforcement*. Imagine precisely this pattern of behavior: Each party is obligated to do all favors that she has the opportunity to perform. If any party fails to do any favor, all parties revert to autarky for the rest of time.

This is a perfect equilibrium, because each party i, called upon to do a favor, compares the immediate cost of doing the favor against the equilibrium continuation value of *all* her relationships; if she fails to do the favor, *no one* will do any favors for her (nor will she do any favors for anyone), for a continuation value of 0. Therefore, this is an equilibrium as long as

$$\max_{j \neq i} m_{ij} \leq \frac{\sum_{j \neq i} \left[\lambda_{ji} a_{ji} - \lambda_{ij} b_{ij}\right]}{r} \quad \text{for each } i; \tag{3}$$

that is, as long as the largest-cost favor that each i is called upon to do is less than her equilibrium continuation value from all relationships, a condition that holds in both examples.[6]

A problem with the "all-favors" equilibrium is that, from the perspective of u-efficiency, favors whose cost exceeds their benefit should not be done.[7] For instance, suppose that the cost–benefit vector (x_{ij}, y_{ij}) has the following probability distribution (for all i and j): With probability 0.8, $(x_{ij}, y_{ij}) = (2, 3)$; with probability 0.1, $(x_{ij}, y_{ij}) = (1, 4)$; and with probability 0.1, $(x_{ij}, y_{ij}) = (5, 4)$. The average values are 2.2 for x_{ij} and 3.2 for y_{ij} and so, in the two examples (that is, with the λ_{ij} and r as in those examples), bilateral reciprocity fails,[8] while all-favors, enforced by social punishment (autarky for all if anyone fails to do any favor), is an equilibrium, with expected payoffs of 10 for each player.

But in the all-favors equilibrium, favors of the third type, with cost–benefit vector $(5, 4)$, must be done, and they are u-inefficient. If we could find a way to do only favors of the first two types, the expected payoff for each party (with $r = 0.1$ and I and the λ_{ij} as in either of the two examples) would be 11. But if we depend on social enforcement, and if social enforcement is based on the entire community being able to verify that each player is acting in the "right" manner, because the information about costs and benefits is local, this information seemingly cannot be used, the classic rules-versus-discretion dilemma.

[6]In this equilibrium, if some i fails to do a favor for some j, *all* favors cease. It might seem more intuitive to punish i by, say, having all favors for her cease. But, at least in some cases, this is problematic. To sustain favor-giving by others, it may be required that i continue to do favors. This is clearest, perhaps, in the second example, where $i+1$ loses any incentive to do favors for $i+2$ if i ceases to do favors for $i+1$. But it also happens in example 1, when and if we reach an (off-path) subgame in which only a few players are left doing favors. It may be possible to punish i by withholding favors from her for some length of time, but this should not be so long that she loses incentive to do favors for others, at least in the simpler forms of rule-based equilibrium.

[7]Pareto efficiency *per se* does not imply that all favors for which $x_{ij} > y_{ij}$ should go undone; e.g., in the efficient outcome that most favors some specific j', all favors for j' should be done.

[8]Both (1) and (2) fail; no bilateral equilibrium except for autarky is possible.

5 Equilibria with Unilateral Discretion

This is not entirely correct. It is possible to punish a player who fails to do a favor sufficiently to keep each player willing to do the low-cost favors, while punishing her in a way and to a degree that keeps her doing favors when punished and somewhat compensates others.

Suppose the distribution of cost–benefit vectors (for all pairs i and j) is the three-outcome distribution just given. Combine this with $r = 0.1$ and the λ_{ij} as in either of the examples from Section 4. Consider the following strategy profiles:

- Each player at each time t is either in state G (for grace) or P (for purgatory), and the actions chosen by players at any point in time depend on the vector describing the state of each player. The game begins with all players in state G.
  - A player in state G does all favors that come along and that cost her 1 or 2. She refuses to do favors that cost her 5.
  - A player in state P does all favors that come her way.
  - If a player in state P fails to do any favor, everyone immediately moves to autarkic behavior, where no favors are done.
  - If a player in state G fails to do any favor, the player is immediately placed in state P and remains there forever.

This is a perfect equilibrium strategy profile:

1. A player in state P is required to do all favors; the most expensive favor she is required to do has cost 5. Favors that are done for this player depend on which of her fellow players are in state G and which in P, but in either case, at least the favors that cost 1 and 2 are done for her, which generate a flow of benefits to her at a rate of 2.8 units per unit time, or a continuation expected value of 28. (This is an underestimate: fellow players in state P generate expected benefits for her that accrue at a greater rate, so her continuation expected flow of benefits is between 28 and 32.) The ongoing expected flow rate of costs to her is 2.2, for a continuation expected cost of 22. So the expected value for her of the ongoing equilibrium (once she is in state P) is at least 6, a bit more than enough to motivate her to do even the most expensive favor.[9]

[9] Even if this player is the first to arrive in state P, she knows that her fellow players will arrive there with probability one, so her equilibrium continuation value is in fact something more than 6. The value 6 is just a lower bound that suffices to show that this is an equilibrium.

2. A player in state G must do favors of cost 1 and 2 but not those of cost 5. Note that her own decision whether to do a favor (when in state G) has no impact on the expected flow of benefits she will receive; it only affects the her future costs. As long as she conforms to the strategy given, one can compute that the equilbrium continuation costs for her are 19.5, so the difference in expected costs of being in state G and state P are $22 - 19.5 = 2.5$, which (when she is in state G) gives her sufficient motivation to do favors of cost 1 and 2 (but not 5).[10]

We can modify this equilibrium somewhat, to allow players who are sent to state P to return to state G, after a while. We do this with a "publicly observable, random alarm clock,"[11] which goes off at an exponentially distributed time with mean $1/\phi$ for some parameter ϕ; we show in the appendix that, for these numbers, as long as $\phi \leq 0.5$, this gives an equilibrium. In fact, we will show that if r is taken as a parameter (instead of the constant 0.1), one has an equilibrium of this form as long as $r + \phi \leq 0.15$, where the most efficient of these (symmetric) equilibria is where $\phi = 0.15 - r$.

By employing unilateral discretion, then, (symmetric) equilibria can be constructed that improve (in terms of payoffs) on the all-favors outcome, at least in some cases, and even in cases where bilateral reciprocity will not be effective. But can the u-efficient outcome be attained? If the u-efficient outcome involves, for some i and j, some i-for-j favors being done but not others (that is, if, for some i and j, neither no-favors nor all-favors is u-efficient), the answer would seem to be no: To motivate i to do some favors, she must be punished for not doing others, and this punishment, presumably, moves the outcome away from u-efficiency. How do we make this rough intuition precise?

Indeed, for other parameterizations—for instance, where bilateral reciprocity will work—the intuition is wrong; if bilateral reciprocity works and if some i-for-j favors should not be done (to enhance efficiency), j can be trusted to punish i (only) for not doing those favors that i should do; inefficient punishment is unnecessary.

Our interest is in situations where bilateral reciprocity cannot work and where, in consequence, community enforcement is required. To develop results that deal with this sort of situation, it is easiest to restrict the strategy profiles (and, therefore, equilibria) at which we look.

Definition 2. An strategy profile is *social* if it consists of (behavior) strategies for each player in which the actions taken by each player at any point in time depend

[10]Detailed computations concerning this equilibrium are given in the appendix.

[11]See the start of the appendix for details about this sort of public randomizing device.

only on information that is either part of the common-knowledge history of play or is "current." More precisely, we will look at various game forms in which actions only take place at the moment that a favor opportunity occurs and where, moreover, only the i, j pair relevant to the particular favor act. In formulating their actions at the moment of a particular i-for-j favor opportunity, this i and j can condition on the public or common-knowledge history of the game and the cost and benefit levels of the current favor opportunity. We use the phrase *perfect social equilibrium* as shorthand for *a perfect equilibrium involving a social strategy profile.*[12]

Note that strategy profiles that employ bilateral reciprocity will not in general be social; bilateral reciprocity in general involves strategies where, if i fails to do a favor for j that, based on its benefit and cost, i was meant to do, j punishes i by withholding later favors. As long as i and j are not alone, and absent any credible broadcast by i and/or j about the cost–benefit vector, the particular values for the cost and benefit do not become public information. Hence, in a social strategy profile, j cannot (subsequently) condition his treatment of i on her earlier refusal to do a favor for him.[13]

It hardly needs saying that restricting attention to perfect social equilibria is not ideal. But this sort of restriction of attention has a long history in the literature, because it provides us with the following result (the proof of which is obvious).

Lemma 1. *In any social equilibrium, the continuation payoffs for all players following the immediate interaction of a pair engaged in a favor opportunity depend only on the public (or common knowledge) history of the game.*

Next section, we will allow players to engage in cheap-talk communication, cheap talk that will trigger future actions and, therefore, continuation values. But if we do not allow this—if the only permitted actions are whether or not to do a favor

[12]Ghosh and Ray (1996) use the same name for a very different object.

[13]Once we make this restriction, we can be more formal concerning what perfection requires. In a social equilibrium, we worry about behavior involving a pair i and j at an instant at which an i-for-j favor opportunity arises—we never will consider situations in which anyone other than i and j act at that instant—and about how the game progresses after the dust of this sort of interaction clears. In a social equilibrium, actions subsequent to a particular favor interaction are based on a history of the game that is common-knowledge to all players; this isn't quite subgame perfection in the formal sense, but it is effectively the same. And within a particular interaction, we assume that i and j are on a common-knowledge basis, at least insofar as their actions are concerned: the past affects current actions only through the common-knowledge elements of their history, while the current information of costs and benefits of the immediate favor are, by assumption, known to both.

when the opportunity arises, and all the universe of players see is (a) that a favor opportunity occurs, where some specific i has the opportunity to do a favor for a specific j, followed by (b) whether i does the favor or not—then some measure of u-inefficiency is unavoidable in social equilibria.

Proposition 1. *If there are more than two players, and if for some pair i and j, there is positive probability that $x_{ij} > y_{ij}$ and positive probability that $x_{ij} < y_{ij}$, then no perfect social equilibrium can give the u-efficient payoffs.*

Proof. As long as there are more than two players, the c.k. information does not include the costs and benefits of any favor. Hence, in any social equilibrium, if the opportunity arises for i to do a favor for j, there can only be two continuation values for i, a continuation value v_1 if she does the favor, and a continuation value v_2 if she does not.

Moreover, the only way to achieve the u-efficient payoffs is with the u-efficient outcome (almost surely), which means that, along the path of play, the favors that are done are precisely those whose cost is less or equal to the benefit generated. So, for the pair i and j such that there are (x_{ij}, y_{ij}) such that $x_{ij} < y_{ij}$ (a favor that is done) and $x_{ij} > y_{ij}$ (a favor that is not done), the two continuation values v_1 and v_2 (both of which are along the path of play) must be identical, both being the value if, subsequently, the u-efficient outcome takes place. But then when $x_{ij} < y_{ij}$, player i will not do the favor: Her continuation value doesn't depend on whether she does it or not (this is a social equilibrium, so j cannot use the information he has to punish her) and, of course, the favor is costly to do.

While we cannot reach u-efficiency in a perfect social equilibrium (under the conditions of Proposition 1) for any fixed $r > 0$, we can *sometimes* approach u-efficiency asymptotically, as r approaches 0. The italicized *sometimes* has to do with the following condition:

(Strict) Condition A. *The data of the problem are said to satisfy Condition A if, for each pair i and j such that $\lambda_{ij} > 0$, if (x, y) is a cost–benefit vector in the support of the distribution of such vectors for i and j, and if $x > y$, then $x \geq M_{ij}$. And the data are said to satisfy the Strict Condition A if such x satisfy $x > M_{ij}$.*

The purpose of this condition and the explanation of *sometimes* can be explained by comparing the following two probability distribution for cost–benefit vectors:
Distribution 1: $(x, y) = (2, 3)$ with probability 0.8, $(1, 4)$ with probability 0.1, and $(5, 4)$ with probability 0.1.
Distribution 2: $(x, y) = (2, 3)$ with probability 0.8, $(3, 2)$ with probability 0.1, and $(4, 5)$ with probability 0.1.

If Distribution 2 is the distribution of cost–benefit vectors for any pair i and j such that $\lambda_{ij} > 0$, then the data fail to satisfy Condition A; u-inefficient pair, namely $(3, 2)$, has cost level 3 that is less than the largest cost level that is part of a u-efficient pair, the 4 in (4, 5). But if, say, the distribution of cost–benefit vectors for all i and j such that $\lambda_{ij} > 0$ look like Distribution 1, the data satisfy Condition A.

The relevance of this distinction arises whenever we allow the favor giver unilateral discretion concerning which favors she performs in a social equilibrium (for more than two players). In a perfect social equilibrium, player i's continuation payoffs can only be influenced by the history of past favor opportunities and those that were done or not. So, insofar as a player has some ability to forgo doing some favors, she will naturally choose not to do those favors that are most costly to herself. These will be the same as the u-inefficient favors (on a pair-of-players-by-pair-of-players basis) if Condition A holds. But if Condition A fails, then some of the i-for-j favors required for u-efficiency will be more costly than some of the favors that must be omitted. And, based solely on her own incentives in any perfect social equilibrium, you cannot have i choosing to do the former when she needn't do the latter. Formally:

Proposition 2. *Suppose the data satisfy Strict Condition A and $\mathcal{A}_i > \mathcal{B}_i$ for each i. Then for sufficiently small r, a perfect social equilibrium exists for the game where payoffs are discounted at the rate r that gives expected payoff vector $(v_1(r), \ldots, v_I(r))$ to the players, where each $rv_i(r)$ approaches $\mathcal{A}_i - \mathcal{B}_i$, the (normalized) payoff associated with u-efficiency, as r approaches 0.*

Conversely, if there are more than two players and the data do not satisfy Condition A, there exists a strictly positive number K such that, if $(v_1, \ldots, v_I)$ is the vector of expected payoffs to the respective players in some perfect social equilibrium for the game with an interest rate r,

$$r\big[v_1 + \ldots + v_I\big] \leq \sum_i \big[\mathcal{A}_i - \mathcal{B}_i\big] - K.$$

In words, without Condition A, the sum of the (normalized) equilibrium value functions is uniformly bounded away from the u-efficient (normalized) sum, in any perfect social equilibrium.[14]

The proof is left to the appendix. (The proof of the positive half mimics standard proofs that employ the strong law of large numbers and give players "budgets" of

[14]This this proposition leaves unaddressed the knife-edge case where, for i and j such that $\lambda a_{ij} > 0$, there is a u-inefficient cost–benefit pair (x, y) such that $x = m_{ij}$. We conjecture that, for this knife-edge case, the second half of the proposition applies. But we are unable to show this.

excused favors. The proof of the second half is more original.) The constant K is identified in the proof.

6 Ex Ante Excuses

The use of unilateral discretion—the favor-giving player unilaterally decides whether to do a particular favor or not—fails to reach efficiency, because each player must have some incentive to perform some favors, hence a favor giver must be "punished" if she fails to do a favor. As the interest rate r goes to zero, punishment (more accurately, what triggers punishment) can be arranged in a manner that, in some cases, gives us asymptotic u-efficiency, an old and fairly well-known story. But in other cases, even a vanishing interest rate fails, because the favor giver is concerned only with her costs (in a perfect social equilibrium), not with the costs and benefits of the favor.

But this fails to use our assumption that the "private" information is local. Can this assumption be employed to solve these two problems? It can, and our objective is to show how this may happen. We do this in two steps: In this section we deal with inefficiencies raised by the need for punishment, for $r > 0$; in Sections 7 and 8, we deal with the problem that, in Section 5 (and in this section), the favor giver is only concerned with her costs.

The idea in this section is relatively simple. Suppose an i-for-j favor arises. Suppose we are looking at an equilibrium in which i will not do this favor, being willing to undergo some form of punishment, instead. The favor receiver j, knowing that this is so, might be willing to "excuse" i with a publicly broadcast speech that "I'm okay with i not doing this favor for me, this time." Since i will not do the favor in any case, j seemingly might as well excuse i.

But it isn't quite as straightforward as that.

In the equilibria of Section 5, to make the efficiency loss as small as possible, we employed a form of punishment for i that improved matters somewhat for j; i went from doing some favors for j to doing them all. If we want j to be willing to excuse i, we can't arrange matters so that j prefers that i is punished. At worst, we want j be neutral to whether i is punished or not. Indeed, to give j a positive incentive to excuse i, we might want j to suffer if he doesn't excuse i, triggering punishment.

Note carefully, in the equilibria of Section 5, punishment happens. Here we are contemplating something quite different: Punishment of i is a threat that keeps i in line, but it is a punishment that is never triggered along the equilibrium path. So the fact that this punishment would hurt j doesn't affect efficiency.

Won't the fact that punishment hurts j mean that i can forego doing any favors for j, knowing that she will be excused by j (since j wishes to avoid the punishment)? The timing here is crucial: j issues the excuse (or doesn't) *before* i must act. So if, in equilibrium, j expects i to do the favor if he doesn't excuse her, he prefers not to excuse her. If he expects that she won't do the favor for him whether she is excused or not, and if her punishment (for not doing the favor when she is not excused) harms him as well, he has a strict incentive to excuse her. We have the following result:

Proposition 3. *If the favor receiver is able to issue cheap-talk, ex ante excuses, if Condition A holds, and if $M_{ij} \leq (\mathcal{A}_i - \mathcal{B}_i)/r$, for all i and j, a perfect social equilibrium exists which implements the u-efficient outcome.*

Proof. The equilibrium is constructed by specifying what is the bad thing that happens if i fails to do a favor when she is not excused, and also what i does if $x > y$ but j fails to excuse i. Let

$$\phi_{ij} = \frac{M_{ij} r}{\mathcal{A}_i - \mathcal{B}_i} ,$$

and suppose that strategies are specified as follows:

1. When an i-for-j favor opportunity arises with cost–benefit vector (x, y), if $x < y$, j does not excuse i and i does the favor.

2. If, on the other hand, $y < x$, j excuses i, i does not do the favor, and all await the next favor opportunity to come along.

3. If $x < y$ and j (by mistake) forgives i, i doesn't do the favor.

4. Regardless of the values of x and y, if j does not excuse i and i does not do the favor, then a publicly observable randomization is conducted where, with probability $1 - \phi_{ij}$, all players "ignore" what just happened, while with probability ϕ_{ij}, everyone moves to autarky. (The inequality assumed in the statement of the proposition ensures that $0 < \phi_{ij} \leq 1$.)

These strategies are clearly social, and they implement the u-efficient outcome. But are they perfect equilibrium strategies? Suppose an i-for-j favor opportunity arises with cost–benefit pair x and y such that $x \leq y$. If j does not excuse i, i is supposed to do the favor, so to get the benefit of this favor, j will not excuse i. And if j does not excuse i, i can do the favor, with an immediate payoff plus continuation value of

$$\frac{\mathcal{A}_i - \mathcal{B}_i}{r} - x,$$

or she can fail to do the favor, with a continuation expected value of

$$(1-\phi_{ij})\left[\frac{\mathcal{A}_i-\mathcal{B}_i}{r}\right]+\phi_{ij}\cdot 0 \quad = \quad \frac{\mathcal{A}_i-\mathcal{B}_i}{r}-\phi_{ij}\left[\frac{\mathcal{A}_i-\mathcal{B}_i}{r}\right] \quad = \quad \frac{\mathcal{A}_i-\mathcal{B}_i}{r}-M_{ij}.$$

Since $x<y$, $x\leq M_{ij}$ by definition, and i is content to do the favor.

On the other hand, suppose that $x>y$. The favor receiver j is supposed to excuse i. If he does, the favor will not be received, but play will continue along the equilibrium path. And if he doesn't excuse i, i will still not do the favor, running the risk of triggering autarky. So excusing i is clearly a best response for j, given i's strategy. As for i, she is supposed not to do the favor whether excused or not. If she is excused, she can forego doing the favor with no adverse consequences (nor does she gain anything by doing the favor), so she won't do it. But what if j fails to excuse her? The same comparison of payoffs as in the case where $x<y$ and i is not forgiven apply, but now, since Condition A holds and $x>y$, it must be that $x>M_{ij}$, so

$$\left[\frac{\mathcal{A}_i-\mathcal{B}_i}{r}\right]-M_{ij}\geq\left[\frac{\mathcal{A}_i-\mathcal{B}_i}{r}\right]-x,$$

where the left-hand side is i's expected payoff if she doesn't do the favor and the right-side side is her payoff if she does. She (weakly) prefers not to do the favor. ■

The key to this construction, of course, is to set the punishment if an unexcused favor is not done to be at a level at which i will do the unexcused favors that are u-efficient, but be willing not to do those that are u-inefficient. Because of the latter and the fact that j suffers from i's punishment, j will excuse i when he is supposed to do so. Of course, different sorts of punishment could be constructed to the same effect; for instance, rather than move to autarky with an appropriate probability, we could arrange an equilibrium in which the omission of an unforgiven favor means a descent into autarky with probability 1, but for an amount of time (either random or deterministic) that is set to give i just the right incentives.

This form of strategy profile can potentially be adopted to any outcome of the form, for each ordered pair i and j, i does all favors for j whose cost is less than or equal to some cut-off value $\hat{x_{ij}}$, and is excused for all favors whose cost exceeds this cutoff. We say "potentially" because, to be an equilibrium, we need that each i is willing to do all the favors she is called upon to do; if v_i is the expected flow of benefits less costs for i in this arrangement, we need that $\hat{x_{ij}}\leq v_i/r$.

Of course, this means that if Condition A fails to hold, this form of strategy profile cannot be used to implement the u-efficient outcome. In fact, we can show that if Condition A fails to hold, no perfect social equilibrium for games where j

speaks before i acts will implement the u-efficient outcome. And, if (1) j is limited to say either "i is excused" or "i is not excused" (or any two-word message space) prior to i acting, and (2) a condition is met that is stronger than failure of Condition A, then normalized perfect social equilibrium payoffs are uniformly bounded away from u-efficiency.

Proposition 4. *Suppose that Condition A fails; that is for some pair i and j with $\lambda_{ij} > 0$, within the support of (x_{ij}, y_{ij}) is a cost–benefit pair $(\hat{x}, \hat{y})$ such that $\hat{y} < \hat{x} < M_{ij}$. Then in any game form in which the favor receiver can issue cheap-talk declarations (only) before i acts, no perfect social equilibrium can implement the u-efficient outcome, for any interest rate r.*

Moreover, suppose (1) the favor receiver is limited to his choice of two (cheap-talk) messages before i acts, and (2) there are four cost-benefit pairs, $(x(k), y(k))$ for this i and j with

$$y(1) > x(1) > x(2) > y(2) > y(3) > x(3) > x(4) > y(4).$$

Then there exists a positive constant K such that, if $(v_1, \ldots, v_I)$ is a vector of expected payoffs to the players in any perfect social equilibrium for interest rate r,

$$r \sum_i v_i \leq \sum_i (\mathcal{A}_i - \mathcal{B}_i) - K.$$

That is, the sum of normalized equilibrium payoffs in any perfect social equilibrium is uniformly bounded away from the sum of normalized u-efficient payoffs. [15]

Proof. We will only show the first part of the proposition here; we prove the uniform bound in the appendix. Suppose by way of contradiction that we could realize the u-efficient outcome in this circumstance. Since $\hat{y} < \hat{x} < M_{ij}$ for some pair i and j, there is a cost–benefit vector $(\check{x}, \check{y})$ for this (ordered) pair in the support of their possible cost–benefit vectors such that $\hat{x} < \check{x} < \check{y}$. Suppose that the u-efficient outcome could be attained (for some interest rate r) in a perfect social equilibrium. Then, in this equilibrium, along the path of play we will reach (with probability 1) a point at which an i-for-j favor opportunity occurs. There is positive probability that the cost-benefit vector for this favor is $(\check{x}, \check{y})$, in which case this favor must be done: Along the path of play (at this point), i will do the favor. j also has some behavior prescribed by the equilibrium at this point. Whatever it is (even if

[15] We emphasize that this part of the proposition depends on the assumption that j can issue only two messages before i decides whether to do the favor. See the discussion following the proof in the appendix for more on this point.

it involves randomization), doing this leads i to choose to do the favor. So suppose that, at this point, the cost-benefit vector is (instead) $(\hat{x}, \hat{y})$. By doing whatever he is meant to do if the cost-benefit vector is $(\check{x}, \check{y})$, j presents i with a choice for which she *must* choose to do the favor: The cost of the favor $\hat{x}$ is strictly less for her than is $\check{x}$, and (in any perfect social equilibrium) her continuation payoffs (if she does the favor and if she does not, given whatever j has said) are the same, so if she weakly prefers to do the favor when her immediate cost is the larger $\check{x}$, she strictly prefers to do the favor when her immediate cost is the smaller $\hat{x}$. But, for j, his continuation payoffs if she does the favor are fixed (in any social equilibrium that implements the u-efficient outcome exactly), and his immediate payoff from inducing her to do the favor is larger, so he will induce her to do the favor. Which, since $\hat{x} > \hat{y}$, means that this is not implementing the u-efficient outcome. ■

7 Simultaneous Announcements

The analysis of Section 6 makes heavy use of the presumed timing: i and j observe x and y; j excuses i or doesn't; i does the favor or doesn't. But once i and j observe x and y, we could imagine different sequences of actions. And with different sequences of actions, we can achieve u-efficiency without Condition A.

Suppose, for instance, that i and j make simultaneous announcements before i must choose between doing the favor or not. As long as the u-efficient outcome provides enough expected value to all players, meaning that

$$\max_{j \neq i} M_{ij} \leq \frac{\mathcal{A}_i - \mathcal{B}_i}{r} \quad \text{for all } i,$$

the u-efficient outcome can be realized in a perfect social equilibrium. This can be done following the construction in Krishna and Morgan (2004): Whenever a favor opportunity arises, i and j must simultaneously announce either "yes" or "no." If they both announce "yes," then i is required to do the favor, where if she doesn't, autarky reigns. If they both announce "no," then i is not required to do the favor, and whether i does the favor or not, the players all continue to the next favor opportunity. And if their announcements are not the same—if one announces "yes" and the other, "no," then autarky immediately reigns. Finally, when any pair i and j see a cost–benefit vector (x, y), they both announce "yes" if $x \leq y$, and i does the favor, and they both announce "no" if $x > y$, and i does not do the favor.

It is clear that this is a perfect social equilibrium that implements the u-efficient outcome. When an i-for-j favor opportunity presents itself, both i and j are motivated to tell the "truth," because they assume the other party will as well, and if they don't comply, autarky ensues, which is worse than continuing with the u-efficient outcome.

We are not satisfied with this as a solution to the problem, however. For one thing, despite any theoretical appeal, the institution of simultaneous announcements is not one we see employed in practice. Perhaps we don't see this practice because: In such an arrangement, each party has a very powerful threat against the other. Suppose an i-for-j favor oppotunity arises in which $x > y$, so both i and j are meant to announce "no" and the favor is not given. Player j, if he can at all credibly do so, can threaten i that he will announce "yes." If i finds this to be credible, depending on the relative cost of this favor versus the equilibrium continuation payoffs, i will knuckle under and also announce "yes." Or, when $x < y$, if i can credibly theaten to announce "no," j may knuckle under and announce "no" as well. The point is that simultaneous announcements with such drastically different consequences depending on those announcements makes the situation very unstable strategically.[16]

8 Ex Post Forgiveness

So we move to a third possible order of events: when an i-for-j favor opportunity arises, i and j observe x and y; i chooses whether to do the favor; and, if i fails to do the favor, j can issue an *ex post* (cheap-talk, public) statement of forgiveness.

To construct perfect social equilibria in which discretion is exercised by the use of forgiveness, we employ the following argument. In a perfect social equilibrium, favor receiver j will have two continuation values at a given point in time if, at that time, i has the opportunity to do a favor for j and fails to do so: j has a continuation value if he forgives her, and he has a continuation value if he does not. Since the decision whether to do the favor is *fait accompli*, j will do whichever of these (forgive or not) has the higher continuation value. Only if the continuation values are identical will he discriminate between the two responses.

But in an equilibrium where i is meant to do some favors for j but not others, and j is meant to forgive i when she doesn't do the "right" favors, we need j to employ both responses. If i fails to do a favor that she is not supposed to do, we need j to forgive her, without adverse consequences to her. So j's continuation value of forgiveness must be at least as large as for not forgiving i. But if it is strictly

[16]For one thing, this puts great weight on the assumption that i and j see exactly the same information. Imagine instead that i and j each see their own payoff-relevant variable precisely, but j sees x plus some noise, while i sees a noisy signal about y. If they are not absolutely certain what their partner will announce, and if the consequences of a mis-announcement are so grave, what should they do?

larger, then i can safely forbear from doing any favors for j, knowing that j (in a perfect social equilibrium) will forgive her.

This suggests the following form for equilibrium strategy profiles, assuming we wish to implement the u-efficient outcome. When an i-for-j favor opportunity arises with cost–benefit vector (x, y), i and j take note of x and y and, in particular, whether $x < y$ or $x > y$. In the former case, i does the favor. In the later case, i does not do the favor, and j forgives her publicly. And, so that i has the incentive to do favors for which $x < y$, j does not forgive such omissions, and i is punished for this in a way that gives her the incentive to do all such favors. That is, i's continuation value if she fails to do a favor and is not forgiven is less than her continuation value after she does a favor by an amount greater than M_{ij}. Finally, j is indifferent *ex post* between excusing an omitted favor and not, so he will always tell the "truth," forgiving those omitted favors where $y > x$ but not those where $x < y$.

Suppose that we can devise punishment for i that is severe enough for i so that she will do all the favors she is meant to do and that weakly enhances the continuation value for j. We can then use a public randomization between this punishment and autarky (which, we presume, is bad for everyone), mixing with a probability that makes j indifferent between denouncing i or not. Of course, this mixing of the original punishment with autarky makes the net punishment of i worse still for her, so she continues to have the incentive to do the favors she is meant to do. The key, then, is to devise punishment regimes that (1) are severe enough for a prototypical i, that she does the required favors, (2) satisfy the perfection constraint that everyone is willing to carry them out, and (3) makes the prototypical j weakly better off.

We can suggest two ways to do this that, potentially, will implement the u-efficient outcome. We refer to the two as the *mild-punishment specification* and the *stronger-punishment specification*.

Recall that $\mathcal{B}_i$ is the expected flow rate of costs i accrues if she does (only) the u-efficient favors, while $\mathcal{C}_i$ is the expected flow rate of costs for i if she does all favors. Recall that $\mathcal{B}_i \leq \mathcal{C}_i$, with equality only in cases where all the favors i might be called upon to do are u-efficient. Dealing with the case $\mathcal{B}_i = \mathcal{C}_i$ adds complications, so for the time being, we temporarily assume that $\mathcal{B}_i < \mathcal{C}_i$ for all i.

Then, in the *mild-punishment specification*:

1. Somewhat similar to the unilateral-discretion equilibria of Section 5, players are at any point in the game either in state G or state P (or everyone is in autarky). All players begin in state G.

2. Player i in state G does all i-for-j favors whose cost is less than its benefit, 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 immediate cost x exceeds the immediate benefit y, and play continues with no additional consequences.

4. But if a player in state G fails to do a favor whose cost x is less than or equal to its benefit y, 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 a 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 aurtarky). 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 whose cost is less than its benefit, for all j who are in state G; if i is in state P, she does all i-for-j favors, 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 cost x of this favor exceeds the immediate benefit, 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, the cost of which is less or equal to its benefit, 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.

The evident difference between the two specifications is that, in the mild-punishment specification, the player is punished (in state P) by being required to do all favors, although she continues to receive favors just as if she were in state G; on the other side of this coin, a player is penalized for failing to do a "warranted" favor for another player regardless of the state the other player. In the stronger-punishment specification, a player in state P receives no favors unless and until she departs for state G; and no one is penalized for failing to do a favor for a player in state P.

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 5. *For the game in which the favor reciever can issue cheap-talk, ex post forgiveness, 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,*

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

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

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

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

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

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

We mix comments about this result with its proof: The first comment is that these strategy profiles are implementing (for $r > 0$) the exact u-efficient outcome, without imposing Condition A or anything similar. In particular, state P is never reached along the equilibrium path of play by any of the players; unlike the construction of

Section 5, state P here is a threat that keeps the players honest along the path of play, but it is not a threat that ever needs to be carried out.

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 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. (In the language of dynamic programming, we only need to verify that each player's strategy is unimprovable in a single step to show that it is optimal, assuming the other players do not deviate at all.)

So, first, we must verify that, when an i-for-j favor opportunity is not done by an i who is in state G, the favor receiver j is indifferent between forgiving the favor or not doing so. This follows 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. For, once we show this, then a mixture of the two extremes can be found that is indifferent to the middle outcome, which is what is meant by the "appropriate choice of randomizing probabilities."

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 with $x \leq y$ hence with $x \leq \max_{j \neq i} M_{ij}$ and 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,[17] 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 stategy) is at least $\mathcal{A}_i - \mathcal{C}_i$ for as long as she is in state P, and it improves to at least $\mathcal{A}_i - \mathcal{B}_i$ once she escapes. (In the mild-punishment regime, i receives at least all her u-efficient favors—she might receive more if others are in state P for a while—and she must do all favors while in P and the u-efficient favors 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

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

which is easily shown to be $[r(\mathcal{A}_i - \mathcal{C}_i) + \phi_i(\mathcal{A}_i - \mathcal{B}_i)]/[r(r + \phi_i)]$. 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 $(\mathcal{A}_i - \mathcal{C}_i)$ 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 u-efficient favors, it suffices to show that her continuation value in state G exceeds her continuation value in state P by at least $\max_{j \neq i} M_{ij}$. 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

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

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

[17] To verify perfection, we must consider all possible starting states.

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 $\mathcal{A}_i/(r+\phi_i)$. 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 the u-efficient favors 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) $\mathcal{A}_i$ 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 $\mathcal{B}_i = \mathcal{C}_i$ for some i? For such an i, all favors that she may be asked to do are u-efficient, so a simple do-all-favors rule will work. 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 $\mathcal{B}_i = \mathcal{C}_i$ is indicative of a further point: We have described two possible perfect social equilibrium regimes, but there are many other candidates. In particular, the 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 malfeasor can make up for her malfeasance, in a way that weakly improves the situation for others. But then motivating players to perform in purgatory means that the equilibrium value needs to 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.

Finally, as long as the u-efficient outcome is better than autarky for each player, either regime can be used to implement u-efficiency for small r.

Corollary to Proposition 5. Suppose that $\mathcal{A}_i > \mathcal{B}_i$ for all i. Then for all small enough r, both strategy profiles (for appropriate ϕ_i) constitute a perfect social equilibrium.

The corollary doesn't explicitly assume that $\mathcal{C}_i > \mathcal{B}_i$ for all i. If $\mathcal{C}_i = \mathcal{B}_i$ 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 $r/\phi_i \to 0$ and $\phi_i \to 0$, both as $r \to 0$.) For those i such that $\mathcal{C}_i > B_i$, the right-hand side of (WP-ICG) goes to infinity as $r \to 0$, so for all small r, this inequality holds. And since $\mathcal{A}_i > \mathcal{B}_i \geq 0$ for all i, the right-hand side of (SP-ICG) goes to infinity as $r \to 0$; this inequality holds for small r. And on the right-hand side of the two ICP inequalities, as $r \to 0$, both $\big[r(\mathcal{A}_i - \mathcal{C}_i) + \phi_i(\mathcal{A}_i - B_i)\big]/\big[r + \phi_i\big]$ and $\big[r(-\mathcal{C}_i) + \phi_i(\mathcal{A}_i - B_i)\big]/\big[r + \phi_i\big]$ approach $\mathcal{A}_i - \mathcal{B}_i > 0$, 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.

9 Imperfect Local Information

In our analysis of local information, 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.[18] 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 information except as is provided via communication between them (which others may share).

In this setting, Proposition 5 is robust to less than perfect information in at least the following sense:

Proposition 6. *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 $1-\epsilon$. Then, if either the first two or the second two inequalities in Proposition 5 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 from Proposition 5, adapted so that players act according to their $\hat{y}$ or $\hat{x}$ "predictions" and with the randomizing probability used to render the favor receiver indifferent suitably adapted, give a perfect social equilibrium with payoffs that are close to u-efficient. And as $\epsilon \to 0$, the equilibrium payoffs approach those of the u-efficient outcome.*

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

[18]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.

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 $\delta \to 0$, 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 sufficient inequalities from Proposition 5 hold strictly, all actions by the players (except for j telling the "truth") are strictly optimal 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 action that is finely balanced in the equilibrium is j's declaration of forgiveness (or not), if i doesn't do a favor, 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 players will go to state P (or, everyone will go to autarky) with positive probability; but as $\epsilon \to 0$, 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 nonneglible impact on equilibrium payoffs.

But 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.[19] 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 7. *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, the u-efficient outcome generates an expected overall payoff greater than 0, or $\mathcal{A}_i > \mathcal{B}_i$ for all i. Then as $r \to 0$, perfect social equilibria can be constructed (for a game that is augmented by cheap-talk announcements) whose normalized payoffs approach the u-efficient payoffs for the players.*

[19] 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.

This result is more or less a direct corollary of the general results of Fudenberg, Levine, and Maskin (1994, Section 8).[20] 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 7, combined with the second half of Proposition 2, 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 the second half of Proposition 2, if the players cannot communicate and if Condition A fails to hold, then for any $r > 0$, normalized equilibrium payoffs are uniformly bounded away from the u-efficient levels. Proposition 7 says that, for small enough (but still strictly positive) r, even if Condition A fails, that uniform bound is violated when the players can communicate. 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)? Given the already considerable length of this paper, we won't chase this down. Instead, we conclude with some easy answers to the question, For a given $r > 0$, when can/cannot the players achieve precisely the u-efficient outcome in a perfect social equilibrium, with less than full knowledge of x and y by both the favor giver and receiver?

In answering this question, we are fairly agnostic as to the rules of the interaction between and 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

[20] We say "more or less" because in our setting, the timing of actions is random. But this poses no serious problems.

involve only cheap talk.

Proposition 8. *Suppose that, for every i-for-j favor that comes along, both i and j have enough information (individually) to distinguish between the events $\{x < y\}$ and $\{x > y\}$. (That is, for each i and j, the conditional probability each assesses based on their respective information endowments for one of these events is 0 and the other is 1, almost surely.) Then as long as either set of inequalities in Proposition 5 hold, the corresponding description of behavior (where j can issue ex post forgiveness) is a perfect social equilibrium that implements the u-efficient outcome.*

The proof is straightforward. The strategies the players employ in the two specifications of behavior depend only on which of these events holds and, as a general proposition, if action A is known by a player to be weakly better than B for every subevent of a given event, then she knows that A is weakly better than B on the larger event. As for the antecedent condition, think, for instance, of a case where the support of (x, y) is such that for each possible value of x there is a unique possible value of y, and vice versa. Then if i knows x, she knows y; if j knows y, he knows x.

A partial converse to this result is:

Proposition 9. *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 $\{x < y\}$ and $\{x > y\}$. And suppose that, in this event, the other player has superior information to the first. Then, regardless of any arrangement of cheap talk, the u-efficient outcome cannot be implemented in a perfect social equilibrium for any $r > 0$.*[21]

Proof. Suppose the u-efficient outcome 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 $\{x < y\}$ and $\{y < x\}$, at which event j has superior information. Because j has superior information, he knows that i assigns positive probability to these events. To implement the u-efficient outcome, it must be that the favor is done if $\{x < y\}$ and not done if $\{x > y\}$, and the continuation from these separate immediate actions is the continuation of the u-efficient outcome. Since j's information is superior to i's, for this to happen, j must know (be endowed with) whether $\{x < y\}$ or $\{x > y\}$. And if j knows that $\{x > y\}$, he can act as if he knew instead that $\{x < y\}$, have the favor done for him, and continue with the u-efficient outcome, which is better

[21] This does *not* say that we cannot approach the u-efficient payoffs asymptotically, as $r \to 0$. This concerns getting to u-efficiency exactly, for fixed $r > 0$.

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 $\{x < y\}$ and $\{y < x\}$, and i who has the superior information, than the argument just given can be refashioned to imply that i knows which of the two situations prevails and knows that j doesn't know this; when $\{x < y\}$, i can act as if she instead knew that $\{y < x\}$, 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 u-efficient outcome is not implemented, which is a contradiction.

The first antecedent assumption in Proposition 9, that some circumstance arises in which one of the players is (conditionally) uncertain whether $\{x < y\}$ or $\{y < x\}$, is the complementary condition to the antecedent assumption in Proposition 8. To that degree, the two are converse. But, of course, the second part of the antecedent assumption in Proposition 9—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 manipuated, we might be able to implement u-efficiency. A simple two-player example 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 Section 8.[22] And, regarding Bob-for-Alice favors, Alice must declare her

[22]Since there are only two players, and Alice and Bob both know x and y for these favors, we can simplify from Section 8: 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 aurtarky. But

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 autarchy, 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.

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 9 is, of course, violated.

On the other hand, while the first antecedent assumption of Proposition 9 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.

10 Externalities, Side-Payments, and the Two Combined

The model we have analyzed is quite simple and stylized, and when seeking to take its insights to real-world phenomena, it is helpful to know of elaborations or extensions that don't change the conclusions. Two extensions are both germane and, taken

it doesn't hurt to follow the behavior rules from Section 8 exactly.

one at a time, easily handled. We have assumed that the parties with the local information are the only parties directly concerned with the decision taken. And we have assumed that these two parties cannot engage in under-the-table negotiations. If we extended the model in either of these two directions, no problems arise. But if we extend in both directions simultaneously, our conclusions no longer hold.

Imagine first that each favor is characterized by a cost x_{ij} that the favor giver i incurs, a direct benefit y_{ij} that the favor receiver j obtains and, *in addition*, external benefits z that accrue to players other than i and j. (For reasons to be explained, we want these to be (weak) benefits to every other player; that is, if z_k is the external benefit to $k \neq i, j$, we want $z_k \geq 0$ and, of course, $\sum_{k \neq i,j} z_k = z$.) We assume that z is known to (at least) i and j.

The decision rule whether to do a particular favor that leads to (and is implied by) u-efficiency becomes $x_{ij} \leq y_{ij} + z$. In principle, then, with the *forgiveness after the fact* cheap-talk regime of Section 8, we can implement in a social equilibrium the u-efficient outcome, as long as the requisite inequalities hold. In practice, the inequality expressing the first condition—that each player is willing to do all favors when in state P—loosens, at least to the extent that players in purgatory continue to benefit from the externalities generated by other players, increasing the continuation value of being in purgatory *vis a vis* the value 0 that arises if autarky breaks out. But the second condition—players are willing to do all u-efficient favors when in state G if the alternative is state P—becomes tighter. The presence of external benefits doesn't affect this inequality directly, since whether in state G or P, a player will continue to receive the benefits generated by parties doing favors for others. But in this setting, more favors (presumably) pass the test $x_{ij} \leq y_{ij} + z$, and the left-hand side of the inequality, which gives the largest cost of a favor that passes this test, therefore becomes larger.

Or imagine that each pair of players i and j have the ability to engage in under-the-table and inefficient sidepayments. If an i-for-j favor opportunity arises, i can secretly propose to j that, if j will forgive (or excuse) i from doing the favor, i will secretly transfer w units of utility to j, at a cost of αw to i, for $\alpha > 1$. We will assume, to keep the story simple, that i makes the proposal, which j must take-or-leave, and if j accepts the proposal, there are no problems in ensuring that the arrangement is carried out: i does transfer the resources to j, and j does the contracted-for bit of excusing or forgiving.

If there are no externalities, the possibility of such under-the-table side payments doesn't affect the u-efficient equilibria of Section 8: The only i-for-j favors that are being done are those whose cost is less than their benefit, so the favor-driven transaction dominates the use of inefficient, under-the-table side payments. (To be precise, even if we assume that j will accept any under-the-table offer that leaves him better off, which is more or less mandated by perfection, the equilibrium of Section 8 affords no offers that i would prefer to doing required favors and that j would accept.)

But now suppose we combine the two extensions: Each favor comes with external benefits z for the other players, and i can make under-the-table, inefficient side-payment offers to j. Then, clearly, the u-efficient equilibrium of Section 8 is in trouble. The favors that, per u-efficiency, should be done are those for which $x < y + z$. But this leaves open the possibility that $y < x/\alpha$, so if i offers to j that she will transfer $y + \epsilon$ to him if he will forgive her for some small $\epsilon > 0$, he is bound to accept, and this costs i only $\alpha(y+\epsilon)$, which, for small enough $\epsilon > 0$, will be less than the cost x of carrying out the favor.

11 Concluding Remarks

Motivated by the observation that discretion is often exercised by permitting exceptions to rules, this paper analyzes settings where this may be observed. It can be argued that the need for a rule arises when self-interested individuals would otherwise opt out of activities that benefit the community as a whole. Then, to enforce the rule, the community would have to observe perfectly the deviation from the rule. When only the deviator has the information required for perfect observation, we arrive at the common rules-versus-discretion trade-off. But when information is not strictly private but instead is local, exceptions to the rule that are "blessed" by a second informed party can be part of an overall equlibrium, rationalizing the "exceptions" we see empirically.

This is not a hugely profound insight, especially when one considers so-called trilateral governance, where a seemingly neutral third party is employed to "learn the facts" and then adjudicate disputes. But in our stylized model, it is the other party most closely affected by the exception that shares the local information and does the required "blessing," and still we see local information employed to enhance efficiency. And we have seen how the timing by which the exception is permitted—via ex ante excuses or ex post forgiveness—can have a major impact on what can be achieved and, in particular, on the incentives that must be created for the second party.

Indeed, since in our analysis it is the second party directly involved who "verifies" for the community that the exception was appropriate, one can wonder why we see so many instances of trilateral governance: Why employ a third party when the two parties directly invoved can achieve what is needed on their own? The answer may lie, in part, in our brief analysis of externalities together with the possibility of under-the-table side payments. The two together are clearly problematic for efficiency, in the sort of equilibria we describe. So when there are externalities and side payments together, the case for employing a (more or less) incorruptible third party is strong, especially where this third party relies on her reputation for incorruptibility to continue to play her economic role. (Of course, this is only part of the case for trilateral authorities; a simple enforcement-of-agreements story works in cases where there are no externalities to worry about, but where even "community enforcement" is inadequate and contractual guarantees are required.)

In our analysis, the information endowments of the players are given exogenously. But it isn't much of a leap from our models to the conclusion that parties involved in these types of transactions would want to invest in and then employ local information. That is, our results informally make the case that efficiency can be enhanced by turning private information into local information and then, via something approximating cheap talk, into globally held information. This "investment in information" is, of course, something we see in the real world: Countries invest in embassies and diplomacy, trade missions, and multi-national trade forums, all of which serve these goals (and, of course, other goals.) As mentioned in the introduction, Toyota spends significant resources to understand suppliers' cost structures and to have suppliers understand the manufacturing environment of Toyota, as well as on facilitating communication among its suppliers. Deans, department chairs, and (more generally) managers of all stripes are encouraged to "get to know" the personal concerns of the people they manage.

Further, when we do observe the use of bilateral discretion it is often in the form of ex-post forgiveness. In the GATT, a violation of the agreement is observed and then trading partners decide if a punishment for the violation should be pursued or if the violation should be excused. Of 369 WTO dispute cases 136 reached full panel process, meaning 233 parties settled "out of court" after the violation had been committed and the matter brought to the attention of the WTO.

We have conducted our analysis with the u-efficient outcome (either exactly or asymptotically) as the desired state of affairs. As we noted, this is only one point on the efficient frontier. But the reader has no doubt recognized that, to the extent that some other efficient outcome can be characterized in terms of a set of favors (in terms of the identity of the two parties involved and the cost and benefit levels of the favor) to be done, with the complementary set not to be done, most of our constructions and results go through in terms of implementing that outcome, either exactly or asymptotically.

Graduate School of Business, Stanford University

12 Appendix

Random alarm clocks

In a number of the constructions of the paper, 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 straighforward 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 $k = 1, \ldots$, a publicly observable randomization is conducted with outcomes "ring" having probability $\phi\delta$ and "not yet" with probability $1 - \phi\delta$. The first occurance 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.

Computations for the equilibria with unilateral discretion

Within the context of the basic model, suppose player i behaves as follows: She begins in state G, in which state she does a selection of favors with 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 a favor that she refuses to do arrives, and she does in fact refuse to do this favor, she moves to state P, in which she does all favors (combined arrival rate $\lambda + \lambda'$). Favors of the second type have an expected cost (conditional on being such a favor) of c'. She stays in state P for an exponential length of time 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.

She discounts costs at an instantaneous rate of r (that is, the present value of cost

x incurred at time t is $e^{-rt}x$); we wish to compute the expectation of the discounted sum of her costs, starting from each of the two stages.

Call the expectation of her discounted costs starting in G, C, 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

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

$$= \frac{\lambda}{\lambda+\lambda'+r}(c+C) + \frac{\lambda'}{\lambda+\lambda'+r}C'. \tag{1}$$

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

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

$$= \left[\frac{\lambda(c+C') + \lambda'(c'+C') + \phi C}{\lambda+\lambda'+\phi+r}\right]. \tag{2}$$

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.

We can apply these formulae to the numerical example that spans Sections 4 and 5. 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. As noted in the discussion in the text, players have no influence over the benefits they accrue from favors done for them, and a lower bound on the expected flow rate of benefits is given by the expected flow rate when all other players are in state G, which is 2.8 units per unit time, for an expected flow rate 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 $C' = 22$. So, with an alternative of triggering autarky, she is willing to do any favor with immediate cost 6 or less. (We did this computation in the text, and it is repeated here only for purposes of completeness.)

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. To reiterate, her decision has no impact on her expected discounted benefits from favors done for her, so 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.)

Note that in the strategy profile we have been investigating, "punishment" 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 $I \times 2^I$.[23] Because of the complexities involved, but even more because it is not the main line we wish to pursue in this paper, we do not pursue the computation of equilibria with unilateral discretion and this sort of punishment.

Concerning the proof of the first half of Proposition 2

The proof of the first half of the proposition uses techniques that are fairly standard, although they must be adapted to this continuous-time, ansynchronous-event setting. 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. Because the methods are standard, we only outline the way we would proceed in this half.

Throughtout this portion of the proof, two assumptions are maintained: (1) Strict Condition A holds, which is that for every ordered pair i and j, if (x, y) is in the support of cost–benefit vectors of i-for-j favors, and if $x > y$, then $x > M_{ij}$. (2) For each player i, if all favors involving i are done when the cost of the favor is less than its benefit, then i, in expectation, has a strictly positive expected value, or $\mathcal{A}_i > \mathcal{B}_i$.

Then the method of proof is, essentially, to construct equilibria that, in the limit, as $r \to 0$, realize the u-efficient (normalized) payoffs. 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 $\psi_{ij} := \text{Prob}[x > y]$, 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 whose cost exceed their benefit 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.

[23] 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.

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.

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}. That is, in each ordered pair i and j, i will omit the most expensive favors for j. 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.

And, because Strict Condition A holds, when she and her peers do this, this means that over each period of length T, the favors that are done (on average, except for an event of small probability) are the u-efficient favors. This means that the expected overall payoff for each players over each of these time intervals is strictly positive; this is where the second maintained hypothesis 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.

Proofs of the second halves of Propositions 2 and 4

The proofs of the second halves of Propositions 2 and 4 use the same technique, so we give them together. In both, we want to show that the sum of the payoffs in any perfect social equilibrium (for a fixed value of r), when normalized, is uniformly bounded away from the normalized sum of the u-efficient payoffs. So fix some r and a perfect social equilibrium. The normalized sum of u-efficient payoffs is

$$\sum_i \Bigg[\sum_{j\neq i} \lambda_{ji} A_{ji} - \lambda_{ij} B_{ij}\Bigg] = r \sum_i \sum_{j\neq i} \sum_{n=1}^{\infty} \mathbf{E}\Big[e^{-rT^n_{ij}}(y^n_{ij} - x^n_{ij})1_{\{x^n_{ij}<y^n_{ij}\}}\Big],$$

where T^n_{ij} is the random arrival time of the nth i-for-j favor and (x^n_{ij}, y^n_{ij}) is the cost-benefit vector for this favor. (If $\lambda a_{ij} = 0$, either omit j from the second sum for this i or take $T^n_{ij} \equiv \infty$.) On the other hand, the normalized sum of payoffs in the equilibrium is

$$r \sum_i \sum_{j\neq i} \sum_{n=1}^{\infty} \mathbf{E}\Big[e^{-rT^n_{ij}}(y^n_{ij} - x^n_{ij})1_{\{i \text{ does this favor for } j\}}\Big],$$

where the expectation being taken involves the exogenously determined random arrival times of favors, the exogenously determined randomly benefit–cost vectors for each favor, and also the endogenously determined strategies being employed in the perfect social equilibrium under investigation. Hence the difference between the normalized sums of the u-efficient payoffs and the equilibrium payoffs is

$$r \sum_i \sum_{j\neq i} \sum_{n=1}^{\infty} \mathbf{E}\Big[e^{-rT^n_{ij}}(y^n_{ij} - x^n_{ij})\big[1_{\{x^n_{ij}<y^n_{ij}\}} - 1_{\{i \text{ does this favor for } j\}}\big]\Big]. \qquad (A1)$$

(We emphasize that the occurrence and timing of favor opportunities is exogenous—that is, unaffected by the strategies the players employ—so writing the difference this way is legitimate.) Note that each term inside the larger square brackets is nonnegative: If $y^n_{ij} > x^n_{ij}$, the first indicator function is 1, so the term inside the smaller square brackets is either 1 or 0; if $y^n_{ij} < x^n_{ij}$, the first indicator function is 0, so the term inside the smaller square brackets is either 0 or -1. Therefore, overall difference is always nonnegative (of course), and we underestimate the difference if we look only at some of the terms in the triple summation.

In the second half of Proposition 2, we assume that for some distinguished ordered pair i^0 and j^0 such that $\lambda_{i^0 j^0} > 0$, there exist $(x(1), y(1))$ and $(x(2), y(2))$ in the support of the distribution of their cost-benefit vectors such that $y(1) > x(1) > x(2) > y(2)$. In Proposition 4, we assume there are four cost-benefit vectors $(x(k), y(k))$ such that $y(1) > x(1) > x(2) > y(2) > y(3) > x(3) > x(4) > y(4)$. In either case, we refer to an i^0-for-j^0 favor with cost-benefit vector $(x(k), y(k))$ as a *favor of type k*, for this ordered pair, and we let p_k be the probability that, in an i^0-for-j^0 favor, the cost-benefit vector makes this a favor of type k. Also, to save on notation, we henceforth drop the superscript 0's on i^0 and j^0; the distinguished ordered pair is simply i and j.

Now go back to the difference in $(A1)$, and examine the term for this specific i and j (ordered) pair. That is, we are looking at

$$r\sum_{n=1}^{\infty}\mathbf{E}\Big[e^{-rT_{ij}^n}(y_{ij}^n-x_{ij}^n)\big[1_{\{x_{ij}^n<y_{ij}^n\}}-1_{i\text{ does this favor for }j}\big]\Big]. \qquad (A2)$$

The occurrence and timing of i-for-j favors is independent of the cost-benefit vectors for those favors, so each time there is an i-for-j favor opportunity, there is probability p_k that it is of type k. Note that for favors of types 1 and 3, the cost is less than the benefit, so they should be done to achieve u-efficiency, while favors of types 2 and 4 should not be done.

We now demonstrate the following: *In any perfect social equilibrium, in the context of Proposition 2, every time there is an i-for-j favor opportunity, either i does not do the favor (with probability 1) if it is type 1, or she does it (with probability 1) if it is type 2.* Which of those two conditions prevail may depend on the specifics of the favor opportunity—what has been the (public) history of play up to that point and perhaps (even) the precise time the favor opportunity occurs. But one or the other must be true, and so for this particular favor opportunity, the expectation in the summation for favor opportunity n is at least

$$\mathbf{E}[e^{-rT_{ij}^n}]\ \min\big\{p_1[y(1)-x(1)],\, p_2[x(2)-y(2)]\big\}.$$

Letting K_0 denote the term $\min\{p_1[y(1)-x(1)],\, p_2[x(2)-y(2)]\}$, this gives as lower bound on (A2), hence on (A1), the quantity

$$r\sum_{n=1}^{\infty}E[e^{-rT_{ij}^n}]K_0=\lambda_{ij}K^0,$$

which establishes the second half of Proposition 2.

And, in the context of Proposition 4, we will demonstrate that *in any perfect social equilibrium, at every occurrence of an i-for-j favor, one of the following four conditions must hold: j does not do the favor if it is of type 1, or j does not do the favor if it is of type 3, or j does the favor if it is type 2, or j does the favor if it is of type 4.* Hence, for

$$K^1=\min\big\{p_1[y(1)-x(1)],\, p_2[x(2)-y(2)],\, p_2[y(3)-x(3)],\, p_4[x(4)-y(4)]\big\},$$

we have the second half of Proposition 4 for $K=\lambda a_{ij}K^1$.

It remains to establish the two italicized assertions. Take the context of Proposition 2, first. When a specific i-for-j favor arises, i must decide whether to do the favor or not. Let v denote her continuation payoff if she does the favor and v' her continuation payoff if she does not. Because the equilibrium is in social strategies, her continuation payoff cannot depend on more than this (and, of course, the public history up to this point, and perhaps the moment in time); in particular, it cannot depend on the prevailing cost-benefit vector. If there is positive probability that she does this favor (at this instant) if it is of type 1, then (because the equilibrium is perfect) it must be that $v - x(1) \geq v'$; in words, her continuation value if she does the favor less the cost of doing the favor must be at least as large as the continuation value if she does not. But, then, since $x(2) < x(1)$, $v - x(2) > v'$, and she *must* do the favor if it is type 2. This establishes that either there is zero probability that she will do the favor if it is of type 1, or probability 1 that she will do it if it is of type 2. That is the first of the italicized assertions, proving the second half of Proposition 2.

For Proposition 4, we must take into account that j can issue one of two messages. Since they are cheap talk, we will simply call the messages A and B. We use v_{AF} for i's continuation value if j issues message A and j does the favor, v_{AN} for i's continuation value if j says A and j does not do the favor, and similarly for v_{BF} and v_{BN}. And we use u_{AF}, and so forth, for j's continuation values.

Suppose at some point in the game, with some public history, at some time, an i-for-j favor opportunity arises.

1.′ Perhaps i's strategy at this point is to refuse to do the favor regardless of what j says. If so, then the italicized assertion is true. So suppose this is not true; i will do the favor with positive probability if (at least) the message is A. (The choice of A is, of course, without loss of generality.)

2.′ But then, by the argument given for Proposition 2, i's strategy if the favor is of type 2, or 3, or 4, and j issues message A must be to do the favor with certainty, since $x(4) < x(3) < x(2) < x(1)$, and we know that $v_{AF} - x(1) \geq v_{AN}$.

3.′ And by a similar argument, if i will do the favor with positive probability when j says B and the favor is of type 2 or type 3, then i with do the favor with probability 1 when j says B and the favor is of type 4. But if i will do the favor with probability 1 regardless of the message when the favor is of type 4, then the italicized assertion is true.

4.′ So the only way the italicized assertion could fail to be true is if i will do the favor with certainty for favors of type 2 and 3, when j says A, and i refuses to do the favor with certainty for favors of type 2 and 3, when j says B.

Now look at j's incentives. When the favor is of type 2, if j says A, he gets $y(2)+u_{AF}$, while if he says B, he gets u_{BN}. If $u_{BN} < y(2) + u_{AF}$, then j will always say A when

the favor is type 2, the favor is done, and the italicized assertion is true. But if $u_{BN} \geq y(2) + u_{AF}$, then $u_{BN} > y(3) + u_{AF}$, j will say B when the favor is type 3, a favor of type 3 is not done, and the italicized assertion is true.
The italicized assertion is true, proving the second half of Proposition 4.

It should be clear that the argument just given for Proposition 4 relies heavily on the assumption that the favor receiver is allowed to send only one of two messages; viz., that the favor giver is excused, or that she is not. We conjecture that, if the favor receiver is allowed as many messages as there are (x, y) pairs in the support of their distribution, asymptotic u-efficiency can be achieved with ex ante communication. Suppose momentarily that players can engage in utility transfers. Having seen the cost-benefit vector (x, y), the favor receiver says to the favor giver, "If you do the favor for me, I will pay you x," if $x < y$. If $x > y$, the favor giver makes no offer. The favor giver, then, is indifferent between doing the favor or not when $x < y$, so we can assume she will do it. The favor receiver, of course, will happily pay x for a favor of value y, if $x < y$, but will be unable to offer any payment sufficient to induce i to do the favor, if $x > y$.

We do not have transferable utility. But, following the general techniques of Fudenberg, Levine, and Maskin (1994), when r is close to zero, the "continuation" of the game can be used in lieu of tranferable utility, as long as all players are aware of the "deal" struck between i and j and the appropriate full-dimensionality assumption is met. There are details to be checked, of course, so we only call this a conjecture. But it seems to us to be a fairly safe conjecture.

References

Barth, Mary E. (2008). "Global Financial Reporting: Implications for U.S. Academics." *The Accounting Review*, Vol. 83, 1159-79.

Bowen, T. Renee (2011). "Forbearance in optimal multilateral trade agreements." Mimeo, Stanford Graduate School of Business.

Bratton, William W. (2003). "Enron, Sarbanes-Oxley, and accounting: Rules versus principles versus rents." *Villanova Law Review*, Vol. 48, 1023-57.

Crawford, Vincent P., and Joel Sobel (1982). "Strategic Information Transmission." *Econometrica*, Vol. 50, 1431-51.

Ellison, Glenn (1994). "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching." *Review of Economic Studies*, Vol. 61, 567-88.

Freed, Daniel J. (1992). “Federal sentencing in the wake of guidelines: Unacceptable limits on the discretion of sentencers.” *Yale Law School Faculty Schoolship Series*, Paper 2050, http://digitalcommons.lawyale.edu/fss_papers/2050.

Fudenberg, Drew, David Levine, and Eric Maskin (1994). “The Folk Theorem with Imperfect Public Information.” *Econometrica*, Vol. 62, 997-1039.

Ghosh, Parikshit, and Debraj Ray (1996). “Cooperation in Community Interaction Without Information Flows.” *Review of Economic Studies*, Vol. 63, 491-519.

Hopenhayn, Hugo, and Christine Hauser (2004). “Trading favors: optimal exchange and forgiveness,” 2004 Meeting Papers 125, Society for Economic Dynamics.

Kandori, M. (1992). “Social norms and community enforcement.” *Review of Economic Studies*, Vol. 59, 63-80.

Kessler, Daniel P. (2011). “Evaluating the medical malpractice system and options for reform.” *Journal of Economic Perspectives*, Vol. 25 (2), 93-110.

Kreps, David M. (2004). *Microeconomics for Managers*. New York: W. W. Norton.

Krishna, Vijay, and John Morgan (2004). “The Art of Conversation: Eliciting Information from Experts through Multi-Stage Communication.” *Journal of Economic Theory*, Vol. 117, 147-79.

Lucas, Robert E. (1980). “Rules, discretion, and the role of the economic advisor.” In Stanley Fischer, ed., *Rational Expectations and Economic Policy*, Chicago: University of Chicago Press, 199-210.

Maggi, G. (1999). “The role of multilateral institutions in international trade cooperation.” *American Economic Review*, Vol. 89 (1), 190-214.

Milgrom, Paul, and John Roberts (1983). *Economics, Organizations, and Management*. Englewood Cliffs, NJ: Prentice-Hall.

Mobius, Markus (2001). “Trading Favors.” Working paper.

Modigliani, Franco (1977). “The monetarist controversy, or should we forsake stabilization policy?” *American Economic Review*, (March 1977), 1-19.

Posner, Richard (2007). ”Rules versus discretion—Posner’s comment.” The Becker–Posner Blog, 9/16/07.

Rosenthal, Robert W. (1979). “Sequences of Games with Varying Opponents.”

Econometrica, Vol. 47, 1353-66.

Schipper, Katherine (2003). "Principles-based accounting standards," *Accounting Horizons*, Vol. 17.

Simons, Henry C. (1936). "Rules versus authorities in monetary policy." *Journal of Political Economy*, (February 1936), 1-30.

Sobel, Joel (2011). "Giving and Receiving Advice." Preprint, Department of Economics, University of California at San Diego.

Wolitzky, Alexander (2012). "Communication and Money in Repeated Games on Networks." Preprint, Stanford University and Microsoft Research.