

RULES WITH DISCRETION AND LOCAL INFORMATION*

T. RENEE BOWEN
DAVID M. KREPS
ANDRZEJ SKRZYPACZ

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 verify how the information was used. Instead, actions may need to conform to simple and widely verifiable rules. We study when discretion in the form of exceptions to a simple rule can be implemented, 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*. *JEL* Codes: C73, D82.

I. INTRODUCTION

Economists have long advanced the notion that, in some contexts, decision makers should forbear from using their best judgment 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; Posner 2007).

The argument for the use of discretion typically comes down to (superior) information held by the decision maker. The counterarguments in favor of the application of a rule vary, but among them (and relevant to this article) is 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*

* Anonymous referees and the editors provided comments and suggestions that significantly improved both the content and exposition of the paper. Comments by seminar participants at UC Santa Barbara, Stanford Graduate School of Business, the Harvard–MIT Theory Seminar, the NBER Organization Economics Meeting, the Midwest Economic Theory Workshop, and the NBER Macroeconomics Within and Across Borders Meetings are also gratefully acknowledged, as is the financial support of the Stanford Graduate School of Business.

© The Author(s) 2013. Published by Oxford University Press, on behalf of President and Fellows of Harvard College. All rights reserved. For Permissions, please email: journals.permissions@oup.com

The Quarterly Journal of Economics (2013), 1273–1320. doi:10.1093/qje/qjt013.

Advance Access publication on April 9, 2013.

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 subassemblies (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 by 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 that a legitimate exception was made.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; that is, 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 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

1. Our use of the adjective “local” echoes the way the term is used in Wolitzky (2013). 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.

(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. 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 article 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 article runs as follows.

The basic model is set up in Section II, with preliminary analysis conducted in Section III. 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. But because the larger community of players does not observe the cost and benefit levels of specific favors, if we rely on community enforcement, discretion about which favors are to be done and which to omit may be limited. This is our rules-versus-discretion conflict.

Sections IV and V provide our main results, in which we ask, if we allow the favor receiver to engage in cheap talk, can the rules-versus-discretion conflict be resolved? In Section IV, we allow the favor receiver to speak before the favor giver must

decide whether to do a given favor; the favor receiver can publicly “excuse” the favor giver from doing some (perhaps inefficient) favors. While there are limits to what can be achieved with this sort of *ex ante* cheap talk, we are able to provide a characterization of what can and cannot be achieved. Then, in Section V, 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 *ex post*. We show that anything that can be implemented with *ex ante* excuses can be implemented with *ex post* forgiveness *and more*, because, once the favor has or has not been done, its particular benefits (and costs) are irrelevant to the continuation of the game. Section VI provides a comparison of *ex ante* excuses and *ex post* forgiveness. Roughly, we argue that while more outcomes can be implemented with *ex post* forgiveness, *ex ante* excuses, when effective, have various “practical” advantages.

For most of the article, we restrict attention to “outcomes” that consist (along the equilibrium path of play) of the time-homogeneous implementation of a *selection* out of all possible favors. Section VII relates this restriction to the set of (first-best) efficient outcomes and extends some of our results to mixtures of selections. Section VIII concludes with a discussion of extensions and variations on our basic model.³

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.

I.A. *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 language.

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

3. An Appendix provides the some technical material related to our results, and an Online Appendix provides a number of bells and whistles that are mentioned but not fully discussed in the article.

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 (2013) 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,” 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. A small literature, beginning with Krishna and Morgan (2004), concerns cheap-talk communication by multiple informed parties; see Sobel (2011, Section V) for a more complete bibliography.

We began by invoking the literature on “rules versus discretion,” so connecting our model to this literature in macroeconomics is worthwhile. In this literature (e.g., Kydland and Prescott 1977), monetary authorities would like to announce a particular path of (contingent) action. If this announcement is believed, the general public will respond in a particular and desirable fashion. But once the public has responded in that fashion, the optimal course of action for the monetary authority changes; if the authority has discretion to act in its (new) best interests, it will not carry out the policy on which basis the general public originally acted. The general public, recognizing this, will not

act initially on the supposition that the authority will keep to its announced policy, at least insofar as the authority retains discretion. So the monetary authority may choose instead to “commit” to some rule, giving up its discretion. Barro and Gordon (1983) and Barro (1986) pick up this story, asking, in essence, what force exists to ensure that the authority will follow the rule. They invoke a reputation construction, based on either an infinite-horizon or bounded-horizon-with-incomplete-information formulation: the (long-lived) authority keeps to the rule to preserve a reputation for doing so.

But for any reputation construction to work, where the point of the reputation is to give credibility to promised (or threatened) future actions and so induce favorable current actions from others, it must be possible for the “general public”—the parties for whom this credibility is required—to verify *ex post* that the first party has behaved in accordance with its reputation. It is here that our model enters the story. We imagine that the information needed to verify *ex post* that the first party has conformed to a contingent rule is not widely held, so the contingent rule seemingly cannot be the basis of a viable reputation. Can a second informed party provide for the general public the required verification and, if so, how?

In some institutional settings, the second informed party has as its profession this task of verification. The second informed party, to the greatest extent possible, maintains its “independence” from the decisions of the first party, and so is trustworthy; think for instance of “independent auditors” (and the controversy of auditing firms that also have consulting relationships with the firms they audit). In the model we explore, we look at a very different situation, namely, where the second informed party, if anything, has the *most* at stake in the decision of the first party, which the second party is called on to verify.

It may be helpful here to think of the case of Toyota and its subcontractors. Toyota wishes to maintain a reputation of being “fair” with its subcontractors. If it can credibly commit to behaving in this fashion in the future, it can induce subcontractors to make sunk-cost investments in the relationship while granting Toyota significant discretion, without detailed contractual safeguards. But subcontractors must worry that, having made those sunk-cost investments, Toyota will resort to a hold-up. Toyota, to be credible in its future actions, must put its reputation for being fair on the line with all its subcontractors simultaneously; together

all subcontractors have sufficient bargaining power vis-à-vis Toyota to keep Toyota “in line.” And here our story enters: (how) can Toyota and one of its subcontractors, which share “local information” not available to the community of subcontractors, employ that information to their mutual advantage, without harming the reputation construction that is essential to the entire system?

II. A BASIC STYLIZED MODEL

We work with variations on the following stylized model.

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

Time is continuous, indexed by $t \in [0, \infty)$.⁴ 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 n th 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 $\{(x_{ij}^n, y_{ij}^n); n = 1, 2, \dots\}$, each of these sequences fully independent of all other such sequences and independent of all the arrival times T_{ij}^n . Then, at the time T_{ij}^n of the n th *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 do not assume 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 0, 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

4. 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 on to do more than one favor at any single point in time.

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.

We assume that the only way one player can provide benefits to another is through these favors. In particular, we assume that utility transfers between the players are *not* possible. In Section VII, we briefly reintroduce the possibility of monetary transfers, at which point we discuss this assumption in greater detail.

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. 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.

Given a history of play, each player evaluates that history by 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 this (infinite-horizon) discounted sum.

II.A. Selections of Favors

Our objective is to address the question, which outcomes can be implemented as perfect equilibria in this game, for a given interest rate $r > 0$, especially if we enlist cheap talk on the part of the players? To gain tractability, we specialize in terms of the “outcomes” we consider and in what we mean by “implement.” The first specialization is that we restrict attention to “outcomes” that are the consistent application of a given *selection* of favors.

For each i and j such that $\lambda_{ij} > 0$, let Σ_{ij} denote the set of all subsets of the support of (x_{ij}, y_{ij}) , with S_{ij} a typical element of Σ_{ij} . And let

$$\Sigma := \prod_{\{(i,j): \lambda_{ij} > 0\}} \Sigma_{ij}, \quad \text{with typical element } \mathcal{S} = (\mathcal{S}_{ij})_{\{(i,j): \lambda_{ij} > 0\}}.$$

In words, an S is a *selection* from all possible favor “types,” where one favor type is distinguished from another by who are the giver and receiver, and what are the cost and benefit. The outcomes we consider are those for which, for some selection S , the favors that are done along the path of play are precisely those that are in the selection S , no more and no less.

Three selections S are of particular interest. (1) Let $S_{ij}^0 = \emptyset$ for each i and j , and let S^0 be the corresponding full selection. This is the null or *autarkic* selection. (2) Let S_{ij}^A be the full support of i -for- j favors, and let S^A be the corresponding full selection. This is the *all-favors* selection. (3) Let

$$S_{ij}^U = \{(x, y) \in \text{the support of } (x_{ij}, y_{ij}) : x < y\} \text{ and } S^U = \left(S_{ij}^U \right)_{\{(i,j) : \lambda_{ij} > 0\}}.$$

That is, S^U is the selection in which the favors that are (meant to be) done are those whose cost is less than the benefit they provide; of course, the outcome that results from this selection will maximize the sum of (expected) payoffs of all the players; the superscript U is for *utilitarian*, and we use the term *u-efficient* to refer to both the selection S^U and the outcome (in terms of payoffs) it engenders.⁵

For a fixed selection S , the following notation will be used. For all ordered pairs i and j such that $\lambda_{ij} > 0$, let

$$A_{ij}(S) := \mathbf{E}[y 1_{\{(x,y) \in S_{ij}\}}], \quad B_{ij}(S) := \mathbf{E}[x 1_{\{(x,y) \in S_{ij}\}}], \quad \text{and}$$

$$M_{ij}(S) := \max\{x : (x, y) \in S_{ij}\}.$$

That is, $A_{ij}(S)$ is the expected benefit j derives from a favor opportunity where i is the favor giver and if only favors in the (sub)selection S_{ij} are performed, $B_{ij}(S)$ is the expected cost to i of a i -for- j favor opportunity, if only favors in S_{ij} are done, and $M_{ij}(S)$ is the most costly favor i is called on to do for j , under the selection S . (If i is called on to do no favors for j in S —that is, if $S_{ij} = \emptyset$ —set $M_{ij}(S) = 0$.) Define

$$A_i(S) := \sum_{j \neq i : \lambda_{ji} > 0} \lambda_{ji} A_{ji}(S), \quad B_i(S) := \sum_{j \neq i : \lambda_{ij} > 0} \lambda_{ij} B_{ij}(S), \quad \text{and}$$

$$M_i(S) = \max_{j \neq i : \lambda_{ij} > 0} M_{ij}(S).$$

5. Section VII begins with a more detailed analysis of the connection between selections and efficiency.

That is, $A_i(S)$ is the expected flow rate of benefits accruing to i if favors in the selection S are done for her, $B_i(S)$ is the expected flow rate of costs she accrues, if she does all the favors she is meant to do in the selection S , and $M_i(S)$ is the most expensive favor i is called upon to do under S .

III. TWO EXAMPLES AND SOCIAL ENFORCEMENT

Since every favor has nonnegative cost, it is evident that autarky—that is, the selection S^\emptyset —can always be implemented as a perfect equilibrium, in which each player has a payoff of zero. But what more is possible?

III.A. Bilateral Reciprocity

Suppose that, for a given selection S , and for each ordered pair i and j ($j \neq i$), the following inequality holds:

$$(1) \quad M_{ij}(S) \leq \frac{\lambda_{ji}A_{ji}(S) - \lambda_{ij}B_{ij}(S)}{r}.$$

Then S can be implemented in bilateral fashion: Each i and j agree to do i -for- j favors specified by S_{ij} and j -for- i favors specified by S_{ji} , as long as the other does as well, with the threat of doing no favors for the other if the other ever fails to do a “required” favor. Because of the additive structure of costs and benefits (i.e., no “returns to scale” in the commission of favors), if the relationship between an i and j breaks down, it has no particular impact on other bilateral relationships. Of course, our assumptions on information (both i and j know the cost and benefit of every i -for- j and j -for- i favor) make this implementable.

But what if bilateral arrangements won’t work? Consider the following two examples.

EXAMPLE 1. *The circle, where 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, ..., 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 counterclockwise neighbor. To finish the example, suppose $r = 0.1$ and each (x_{ij}, y_{ij}) pair is degenerate at (2,3). The point is that, clearly, for no pair will

bilateral reciprocity work except for implementing pairwise autarky; for pairs i and j where $j = i + 1$ inequality 1 holds, but it fails when we reverse their roles. (For all other pairs i and j , $\lambda_{ij} = \lambda_{ji} = 0$ —the two never interact—so in that sense, the inequalities are irrelevant.)

EXAMPLE 2. *Star-shaped relationships.* For, say, $I = 21$, suppose that $\lambda_{1j} = \lambda_{j1} = 0.1$ for $j \neq 1$, and $\lambda_{ij} = 0$ if neither i nor j is 1. In other words, everyone (other than 1) can give and receive favors from player 1, but no other pairs interact. Suppose $r = 0.1$, (x_{1j}, y_{1j}) has degenerate distribution at $(9, 10)$, and (x_{j1}, y_{j1}) is degenerate at $(1, 10)$, for each $j \neq 1$. For each $j \neq 1$, the expected present value of the flow of favors less costs (if 1 and j carry out all favors) is $\frac{0.1(10-1)}{0.1} = 9$, more than enough so that j has incentive to keep a bilateral all-favors arrangement with 1. But for player 1, expected present value of favors received from some j less the cost of favors done for this j is $\frac{0.1(10-9)}{0.1} = 1$, which is insufficient to induce her to do favors of immediate cost 9, even if failing to do so means that this j will never do favors for her in the future.

In these examples, if we rely on bilateral reciprocity, no favors will be done. But suppose instead we imagine that players are punished by the whole community if they fail to do a favor for any one member of the community. More precisely, imagine that players do all favors for one another unless and until someone fails to do a favor for someone, at which point *all* players move to autarky, forever. For this to be an equilibrium, it is necessary and sufficient that, for each i ,

$$(2) \quad M_i(S^A) \leq \frac{A_i(S^A) - B_i(S^A)}{r};$$

that is, the most expensive favor that each i is called on to do in the all-favors selection is less than her continuation value of keeping with the prescribed strategy profile. In both examples, inequality 2 is met for each i : in the circle example, i does favors for $i + 1$ so that $i - 1$ will continue to do favors for her; in the star-shaped arrangement, player 1 is willing to do favors costing her 9, to preserve the expected net flow of value she receives from the twenty relationships she has with the various $j = 2, \dots, 21$. Hence, in the two examples, these strategy profiles constitute a perfect equilibrium that implements S^A .

III.B. *But All-Favors May Be Inefficient*

Suppose that, in the context of Example 1 (the circle), cost-benefit vectors for each pair (i, j) (where $j = i + 1$) have the following probability distribution: 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, $(x_{ij}, y_{ij}) = (5, 4)$. For the all-favors selection, we have $A_i(S^A) = 3.2$ and $B_i(S^A) = 2.2$, so all-favors with the threat of autarky for all constitutes a perfect equilibrium, with payoff 10 for each player. But this selection is inefficient: The third type of favor, with cost-benefit vector $(5, 4)$, would be better left undone; if we could find a way to do only favors of the first two types (for the selection S^U), the expected payoff for each party (with the rest of the data taken from the circle example) would be 11.

It is unclear, though, how we might implement this efficient selection as a perfect equilibrium. Because this is based on the circle structure of interactions, bilateral reciprocity necessarily fails. In terms of engaging social enforcement, players other than i and $i + 1$ cannot tell if i is doing the right thing, when an i -for- $i + 1$ favor opportunity arises. Both i and $i + 1$ know the values of x and y in any particular instance, so they both know if this is a favor that ought to be done or not. But for other players, this information is lacking. Insofar as we rely on social enforcement, if the specific favor is not in the selection, i should not be “punished” for failing to do it. But, then, how will the wider community know whether to punish i for an undone favor? If i goes unpunished for not doing a favor, then i has no incentive to do any favor, so some punishment (or, alternatively, some reward for doing a favor) is required.

This is the rules-versus-discretion dilemma within our model. “Do all favors” allows for social enforcement, because everyone can verify ex post whether this rule is being followed. But, as long as $I \geq 3$, the *only* selections that provide each player with enough information to tell whether the selection/rule is being followed are selections S such that, for each i and j , S_{ij} is either \emptyset or S^A_{ij} . To the extent that we rely on social enforcement, it might seem that these are the only “rules” that can be implemented; any finer discretion in choosing which favors to do and which to omit is impossible.

This conclusion overstates matters. Even if we rely on social enforcement, equilibria can be created in which favor giver i unilaterally exercises discretion about which favors to do for a

given j . These equilibria, though, do *not* involve the implementation of a fixed selection of favors in time-homogeneous fashion. If i fails to do a favor for j , in a case where i is meant to do some favors for j , i must be punished. The punishment must be severe enough so that i is willing to do some favors for j , but (at least in examples such as the circle), not so severe that i ceases to do all favors for others, unless the punishment is to be autarky for all.

Analysis of unilaterally exercised discretion takes us too much out of our way; we want to stick to the time-homogeneous implementation of a fixed selection, by involving the favor receiver, who also knows the costs and benefits of a given favor, in the exercise of discretion. So we leave further discussion and analysis of unilateral discretion for concluding remarks and the Online Appendix.

III.C. Social Strategy Profiles and Perfect Social Equilibria

Besides restricting attention to equilibria that implement in time-homogeneous manner a fixed selection of favors, we also restrict attention to equilibria in strategies that satisfy the following formal definition. This definition is somewhat tuned to our specific game: we only allow players to take actions at the moment when a favor opportunity arises, and we only allow the immediate (prospective) favor giver and favor receiver to act. Within this sort of game:

DEFINITION 1. A strategy profile is *social* if it consists of (behavior) strategies for each player in which the actions taken by the prospective favor giver and receiver in their interaction at a given favor opportunity depend only on (a) information that is part of the common-knowledge history of play, (b) the cost and benefit levels of the current favor opportunity, and/or (c) actions taken in this interaction by the other party (and observed by the first party). A *perfect social equilibrium* is a perfect equilibrium in which the players employ a social-strategy profile.⁶

6. Ghosh and Ray (1996) use the same name for a very different object. Please note that a “perfect equilibrium involving a social-strategy profile” is a social-strategy profile in which each player, using a social strategy, is playing an unrestricted best response against the strategies of her opponents. In testing whether we have an equilibrium, we do not require that deviations are social.

This is a serious restriction: consider, for instance, strategy profiles that employ bilateral reciprocity, 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 such strategies are not social.

Why restrict the strategy profiles we are willing to consider in this fashion? Our interest is in cases where enforcement requires social action. It may be possible to construct equilibria in which the continuation play (after some action) is not common knowledge to all the players and yet that involves some level of social enforcement. But we are unable (at least, at present) to shed light on such equilibria. Put more positively, we will show how, by involving the favor receiver in cases where a favor is undone, we can implement some selections S with a perfect social equilibrium. To the extent that having social strategies is a virtue, this sort of result is strengthened by the restriction.

This type of restriction is common in the literature. It appears first (we believe) in Radner, Myerson, and Maskin (1986), and is formalized as *public strategies* and *perfect public equilibrium* in Fudenberg, Levine, and Maskin (1994). Because (once we introduce cheap talk) we examine games in which players can respond at a given favor opportunity to each other's immediate-past actions, it is ambiguous whether our definition is a special case of those in the literature; to avoid confusion, we use a slightly different name. Be that as it may, our concept is clearly closely connected to others in the literature, and it is advanced for the same reason: it provides 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.⁷

7. Once we restrict attention to social strategy profiles, this lemma allows us to 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 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

IV. EX ANTE EXCUSES

We come to the main question: if the information needed to implement a selection S is *local*, meaning held by both the potential favor giver and receiver but not by others, (how) can the potential favor receiver be enlisted to help implement (as a perfect social equilibrium) S ?

We do not change the rule that says that the potential favor giver must decide unilaterally whether to grant the favor. We do not change any payoffs, contingent on the choices favor givers make in that regard. But we allow the favor receiver to issue public, cheap-talk statements, intended to guide the behavior of others.

Of course, cheap talk can take many forms. We could imagine the potential favor giver and receiver both engaging in cheap talk.⁸ Even if we restrict to cheap talk by the receiver only, we could imagine him talking before the favor decision is made and then again after. The language employed could be quite large.

Here we look at the following simple situations. The only party that engages in cheap talk is the favor receiver. In this section, he speaks before the favor decision is made. In the next section, he speaks after the favor decision has been taken. These

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.

8. In the context of this game, where both know precisely the cost and benefit of a given favor, and where there is a form of punishment—namely autarky—that is both dire and, on its own, an equilibrium, simultaneous cheap talk can easily be used to implement a selection S , assuming that Condition A in Proposition 1 holds for S : Prior to the moment when the favor must be done or not, but after both the favor giver and receiver have seen the pair (x, y) , they simultaneously announce “It should be done” or “It should not.” If their messages coincide, what they announce happens. If their messages disagree, autarky immediately ensues, forever. If there is any deviation after a common message (say, they announce “It should be done” and then it is not), autarky immediate ensues. There are well-known strategic issues arising from this sort of equilibrium; for example, each pair of favor giver and receiver, if they can address one another privately prior to the announcement, can threaten the other. There are more subtle issues to resolve if their information is not identical; one then looks for whether, with a richer language, approximate efficiency can be achieved. We will not pursue this further; in this article, the only cheap talker is the favor receiver.

relatively simple possibilities already raise a number of subtle and interesting issues to be explored before (perhaps) going on to more complex patterns and language of cheap talk.

We reiterate a point made in the Introduction: cheap-talk declarations, especially *ex post*, by otherwise disinterested parties are a common phenomenon. (In transaction-cost economics, the term *trilateral governance* is sometimes used to describe such situations.) Auditors, for instance, are meant to verify (after the fact) that financial reports issued by a company are accurate. As a (theoretically) disinterested party, auditors have no incentive to do anything other than to tell the truth.⁹ In our story, however, it will be up to the other party most directly affected by the decision—the favor receiver—who must make the cheap-talk declaration. We find this interesting on grounds that this is the party most likely to share in the information required for efficient arrangements.

With *ex ante* cheap talk, the question *Can selection S can be implemented (for a given r) in a perfect social equilibrium?* has a simple answer.

PROPOSITION 1. If favor receivers are able to issue cheap-talk declarations (only) before the decision whether to do the favor, selection S can be implemented for a given $r > 0$ if and only if

$$(a) \quad M_i(S) \leq \frac{A_i(S) - B_i(S)}{r} \quad \text{for all } i \quad \text{and}$$

$$(b) \quad (x, y) \notin S_{ij} \text{ implies } x \geq M_{ij}(S) \quad \text{for all } i \text{ and } j \text{ such that } \lambda_{ij} > 0.$$

Moreover, if Conditions (a) and (b) hold for S , then S can be implemented with a cheap-talk vocabulary of only two words:

9. Of course, were they completely disinterested, they would have no incentive to say anything in particular; truth-telling is but one equilibrium. Insofar as effort is required to get at the truth, they have no incentive to expend effort, until one enriches the model. Imagine that there is some chance that, after audit, it may come out that the auditor lied or did not show due diligence in the audit procedures. Then a well-compensated auditor, who loses her reputation and ability to earn that compensation if she is caught, has plenty of incentive to (a) work and (b) tell the truth. In the real world of auditing, one has to consider the countervailing incentive for auditors of providing consulting services to clients or, simply, desiring to please a client in order to get the audit engagement again in the future.

A prospective favor receiver either says "the favor giver is excused from doing the favor, this time," or "she is not excused; I expect her to do the favor."

Proof. The necessity of Condition (a) should be clear; this is just the obvious incentive-compatibility constraint along the path of play. But, to spell it out, suppose Condition (a) did not hold. Then there is some ordered pair (i, j) and an i -for- j favor $(x, y) \in S_{ij}$ such that $x > \frac{A_i(S) - B_i(S)}{r}$. If an i -for- j favor opportunity occurs with cost-benefit vector (x, y) (and this happens at some point with probability 1) i 's continuation payoff, if we are implementing S and she does the favor is $\frac{A_i(S) - B_i(S)}{r} - x < 0$. She is assured of no worse than 0 by doing no more favors, hence doing this favor is not a best response by i .

Suppose that Condition (b) does not hold. That is, for some ordered pair (i, j) with $\lambda_{ij} > 0$, there is an i -for- j favor type (x, y) with $x < M_{ij}(S)$. Suppose S could be implemented in a perfect social equilibrium with ex ante cheap talk. Let $(x', y') \in S$ be an i -for- j favor type where $x' = M_{ij}(S)$. When an i -for- j favor opportunity arises (which happens with probability 1 infinitely often), it could be of type (x, y) and it could be of type (x', y') . Since this perfect social equilibrium implements S , there is some cheap-talk statement that j can make, call it "Z", that causes i (in equilibrium) to do the favor with probability 1, if the favor is of type (x', y') . Note that, after this happens, i has the continuation value $V^* = \frac{A_i(S) - B_i(S)}{r}$ and j has the continuation value $U^* := \frac{A_j(S) - B_j(S)}{r}$, since this equilibrium implements S . Let V be i 's continuation value if she does not do the favor; since she is willing to do the favor, it must be that $V^* - x' \geq V$.

Now consider if the i -for- j favor is of type (x, y) . Since the path of play implements S , there is something j says that allows i not to do the favor. After saying this (and not getting the favor), j 's continuation value is, once again, U^* .

So suppose that the favor is of type (x, y) , and j issues the cheap-talk declaration "Z." Player i must choose between doing the favor, for a payoff of $V^* - x$, or not doing it, for a payoff of V . (Here the assumption that the strategies are social comes into play: Continuation payoffs can only depend on what j says and whether or not i does the favor.) Since $x < x'$ and $V^* - x' \geq V$, it follows that $V^* - x > V$, so i will surely do the favor (because the equilibrium is perfect). And, then, j gets the value $y > 0$ of the favor, and same continuation value U^* that he would have had,

had he followed the prescribed equilibrium. Therefore, this was not an equilibrium.

Now assume that Conditions (a) and (b) hold for some fixed \mathcal{S} . Because all favor types have a strictly positive cost, Condition (a) holding implies that either $\mathcal{A}_i(\mathcal{S}) > \mathcal{B}_i(\mathcal{S})$ or i is called on to do no favors in the selection \mathcal{S} . In the latter case, implementation of i 's role is trivial. So we assume that $\mathcal{A}_i(\mathcal{S}) > \mathcal{B}_i(\mathcal{S})$ for all i , leaving the case where this strict inequality holds for only a subset of the i 's to the reader.

We construct a simple cheap-talk regime that implements the selection \mathcal{S} : For each i and j let

$$\phi_{ij} := \frac{M_{ij}(\mathcal{S})r}{\mathcal{A}_i(\mathcal{S}) - \mathcal{B}_i(\mathcal{S})}.$$

Note that per the assumption of the proposition, $0 \leq \phi_{ij} \leq 1$. Specify the following strategies. When an i -for- j favor opportunity arises with cost-benefit vector (x, y) :

- (i) If $(x, y) \in \mathcal{S}_{ij}$, j does not excuse i and i does the favor.
- (ii) On the other hand, if $(x, y) \notin \mathcal{S}_{ij}$, then j excuses i , i does not do the favor, and all await the next favor opportunity to come along. But if j does not excuse i , i even so does *not* do the favor.
- (iii) If $(x, y) \in \mathcal{S}_{ij}$, and j (by mistake) excuses i , i doesn't do the favor. Whether i does the favor or not (whenever she is excused), play simply proceeds without further immediate incident to the next favor opportunity.
- (iv) 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 and move to the next favor opportunity, while with probability ϕ_{ij} , everyone moves to autarky.¹⁰

These strategies are clearly social, and they implement the selection \mathcal{S} along the path of play. We must check whether they are perfect equilibrium strategies. If ever play moves to autarky,

10. It will be clear from the proof that, alternatively, we could in these circumstances have all players move to autarky for a length of time that inflicts the same expected costs on i as does this random descent into autarky. A more interesting issue is that this "punishment" is inflicted not only on i , but as well as on j and on everyone else. We discuss this point after the proof.

continued play is clearly in equilibrium. So suppose an *i*-for-*j* favor opportunity arises with cost-benefit pair $(x, y) \in S_{ij}$. 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

$$\frac{A_i(S) - B_i(S)}{r} - x,$$

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

$$\begin{aligned} (1 - \phi_{ij}) \left[\frac{A_i(S) - B_i(S)}{r} \right] + \phi_{ij} \cdot 0 \\ = \frac{A_i(S) - B_i(S)}{r} - \phi_{ij} \left[\frac{A_i(S) - B_i(S)}{r} \right] \\ = \frac{A_i(S) - B_i(S)}{r} - M_{ij}(S). \end{aligned}$$

Because $(x, y) \in S_{ij}$, $x \leq M_{ij}(S)$ by definition, and *i* is content to do the favor.

On the other hand, suppose that $(x, y) \notin S_{ij}$. 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 forgo 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 $(x, y) \in S_{ij}$ and *i* is not excused apply, but now, since $(x, y) \notin S_{ij}$ implies that $x \geq M_{ij}(S)$ (Condition (b)),

$$\left[\frac{A_i(S) - B_i(S)}{r} \right] - M_{ij}(S) \geq \left[\frac{A_i(S) - B_i(S)}{r} \right] - x,$$

where the left-hand side is *i*'s expected payoff if she doesn't do the favor and the right-hand side is her payoff if she does. She (weakly) prefers not to do the favor, showing that our purported equilibrium is indeed a (perfect) equilibrium. ■

Note that in the equilibrium we construct (if (a) and (b) hold), *everyone* suffers if *i* fails to do a favor for which she is not excused. (Of course, this is an out-of-equilibrium occurrence; it never happens along the path of equilibrium play.) It is unnecessary for the

“innocent bystanders”—everyone except for i and j —to suffer, and depending on the detailed structure in a particular example, it may be possible to arrange matters so that only i and j suffer.¹¹ But to maintain the equilibrium, i ’s punishment cannot be designed in a way that benefits j and, in fact, we suggest that j ought to be punished as well as i : player i ’s punishment is set so that she will do all favors in S_{ij} and, whether excused or not, will do no favors not in S_{ij} . It is then j ’s responsibility to excuse her, if the favor type is not in S_{ij} . If j does not excuse i and i does not do the favor, *either i or j* has failed to follow the equilibrium prescription, and third parties cannot tell who it is. So it is somewhat natural to punish both i and j . Put another way, if j suffers in these circumstances, j has a strict incentive to excuse i when an i -for- j favor type is not in S_{ij} . To the extent that equilibria with strict best responses are more robust, punishing j in these circumstances is a good idea.

V. EX POST FORGIVENESS

We turn now to cheap talk by the prospective favor receiver *after* the favor giver has made her choice whether to do the favor. We have in mind that if a favor goes undone, the (now disappointed) favor receiver can either publicly forgive the favor giver or not.

Ex post cheap talk has one major advantage over ex ante cheap talk. Cheap talk undertaken before the decision whether to do the favor must take into account the cost and benefit levels of the current favor; the favor could still be done, so those values are still germane. But with ex post cheap talk, the “current” favor is now in the past and is irrelevant to continuation values.

On the other hand, ex post forgiveness of the sort we are seeking is delicate in at least one respect. Reason as follows: 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. Because the decision not to do the favor at this point is *fait accompli*, j will do whichever of these (forgive or not) provides

11. But care must be taken here. In, for instance, the circle example, punishing any player in a way that removes her incentives to do favors destroys the possibility of any further favors being done.

him the higher continuation value. Only if the continuation values are identical will he discriminate between the two responses.

But if the selection S requires i to do some favors for j but not others, and j is meant to forgive i when she doesn't do a favor that is not in S_{ij} , 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 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, in implementing a selection S in which some *i-for-j* favors are done and others are not. When an *i-for-j* favor opportunity arises with cost-benefit vector (x, y) , if $(x, y) \in S_{ij}$, then i does the favor. If $(x, y) \notin S_{ij}$, i does not do the favor, and j forgives her publicly. To provide i with the incentive to do favors for which $(x, y) \in S_{ij}$, if i fails to do such a favor, j does not forgive her, and she 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 $M_{ij}(S)$. Finally, j is indifferent ex post between excusing an omitted favor and not, so he will always tell the "truth," forgiving those omitted favors for which $(x, y) \notin S_{ij}$ but not those that are $\in S_{ij}$.

How do we make j indifferent ex post? Suppose the punishment that is devised for i (if she fails to do a favor and is not forgiven) is such that j weakly *benefits* from it. If this is so, then if i fails to do a favor and is not forgiven, we can conduct a publicly observable randomization between punishing i and moving everyone to autarky; as long as autarky is worse for j than not punishing i (which it must be, for individual rationality to hold), a randomizing probability can be selected that exactly balances the continuation values for j if she forgives i and the expected value if she does not and either j is punished or autarky descends.

The key, then, is in devising punishment regimes. There are a number of ways this might be done; all the ones we have been able to construct have fairly complex off-path dynamics, but one is distinguished in that it gives the following remarkable result.

PROPOSITION 2. If favor receivers are able to issue cheap-talk declarations (only) after the decision is taken whether to do the favor, selection S can be implemented for a given $r > 0$ if and only if

$$(a) \quad M_i(S) \leq \frac{A_i(S) - B_i(S)}{r} \quad \text{for all } i.$$

Moreover, the cheap-talk vocabulary required along the path of (equilibrium) play involves favor receivers saying “ i is forgiven, this time” or “no forgiveness,” if i does not do a favor, and the second of these two is never used along the path of play. (Favor receivers are also called on to issue cheap-talk declarations in off-the-path situations, and as with along-the-path declarations, a vocabulary of two possible messages suffices.)

Proof. To take the negative half of the result first, if S is implemented in a perfect social equilibrium, then the payoff to player i is $\frac{A_i(S) - B_i(S)}{r}$. Called on to do a favor from S , i 's continuation payoff must exceed the cost of this favor. So Condition (a) is certainly necessary for the implementation of S .

For the positive half, we construct a perfect social equilibrium that implements S . To avoid special cases in the proof, assume that, for each i , $A_i(S) > B_i(S)$.¹²

To motivate our construction, note that Condition (a) gives us no “slack”: suppose Condition (a) holds with equality for player i . Let j be the corresponding player for whom $M_{ij}(S)$ attains the maximum, and let (x, y) be the i -for- j favor type that has $x = M_{ij}(S) = \frac{A_i(S) - B_i(S)}{r}$. Since we are implementing S , if player i does this favor, as she is meant to do, her continuation payoff will be $\frac{A_i(S) - B_i(S)}{r}$. So if she does not do this favor and is punished, her continuation payoff must be 0. (It can't be less, because 0 is her maxmin. It can't be more, or she won't do this favor.) What we will do, then, is to design punishments for players in which their continuation values (if they are in the midst of punishment) is always 0.

12. If $A_i(S) = B_i(S)$, then Condition (a) implies that i is called on to do no favors in S , so $B_i(S) = 0$, and then $A_i(S) = B_i(S) = 0$ implies that i receives no favors, so i is seemingly a dummy player. For purposes of off-path play, such an i may still have a role to play, so she isn't a dummy player in the full sense of the term; it is those complications we leave to the reader to work through.

We do this with a six-part description of the strategies the players employ:

- (i) Each player at any point in the game will be in one of two states, G (for “grace”) or P (for “purgatory”). All players start in state G.
- (ii) If an i -for- j favor arrives with cost-benefit vector (x, y) when j is in state G, then i does the favor if $(x, y) \in S_{ij}$ but does not do the favor if $(x, y) \notin S_{ij}$.
- (iii) If an i -for- j favor arrives, when i and j are both in state G, and if i does the favor, the game simply continues to the next favor opportunity, with no change in anyone’s status.
- (iv) If an i -for- j favor arrives when i and j are both in state G, and if i does not do the favor, then j issues a cheap-talk declaration: he forgives i if the favor is not in S_{ij} , and he does not forgive her if the favor is in S_{ij} . If j forgives i , the game continues to the next favor opportunity, with no change in anyone’s status. If j does not forgive i , then a publicly observable randomization takes place with one of two outcomes: either i is sent to state P or everyone moves to autarky, forever.
- (v) If an i -for- j favor arrives when j is in state P, i does not do the favor, regardless of the type of favor. Whether i does the favor (by mistake) or not, play continues to the next favor opportunity with no change in status for any player.

So, i ’s punishment (when she has been sent to P) is that no one does any favors for her, unless and until she returns to state G. But we still want her to do favors for other players who are in G, at least those favors that are called for under S . The final part of the strategies concerns how this happens; that is, what takes place if i is in P, j is in G, and an i -for- j favor opportunity arises:

- (vi) If an i -for- j favor arrives when i is in state P and j is in state G, i does the favor if it is in S_{ij} and does not do it if it is not in S_{ij} .
 - (a) If i does not do the favor, play continues to the next favor opportunity, with no change in the status of any player.
 - (b) If the favor is of type $(x, y) \notin S_{ij}$ and i does the favor, j issues a cheap-talk declaration that “ i is not restored to grace.” A publicly observable

randomization is then conducted: either no change in the status of any player occurs, or player j is sent to state P. The probabilities for this randomization will be specified in a bit.

- (c) If the favor is of type $(x, y) \in S_{ij}$ and i does the favor, j issues one of two cheap-talk declarations: Either j says “ i is restored to grace,” in which case i is moved back to state G, and the game proceeds, or j says “ i is not restored to grace,” in which case a publicly observable randomization takes place, with one of two consequences: no one changes status (and we proceed to the next favor opportunity), or j joins i in state P. The exact probabilities for j ’s choice of which cheap-talk proclamation to make and for the publicly observable randomization will be specified in a bit.

We assert that these strategies are social. Specifically, after the dust settles on each favor opportunity, the publicly available information is sufficient to tell the state (G or P) of each player, which then is sufficient to tell what happens at the next favor opportunity.

If players follow these strategies, starting from when all players are in state G, each player i has an expected payoff of $\frac{A_i(S) - B_i(S)}{r}$, which is also her continuation payoff if all are in state G, following any favor opportunity. We define the (as yet unspecified) randomizing probabilities so that each player i , if in state G, is weakly better off the more (in the sense of set inclusion) of her opponents are in state P.

We need notation for the continuation values of each player i , as a function of the set of players (other than i) who are in state P and i ’s state. Let \mathcal{I} denote a generic subset of the set of players less $\{i\}$, including the null set, and write $v_i^G(\mathcal{I})$ for the continuation value for player i , if i is in state G and the players currently in P are given by \mathcal{I} . And let $v_i^P(\mathcal{I})$ be i ’s continuation value, if she is in state P together with the members of \mathcal{I} . Note that the assertion of the last sentence in the previous paragraph is that $v_i^G(\mathcal{I}) \geq v_i^G(\mathcal{J})$ if $\mathcal{J} \subset \mathcal{I}$ so that, in particular, $v_i^G(\mathcal{I}) \geq v_i^G(\emptyset) = \frac{A_i(S) - B_i(S)}{r}$.

To specify the probability in the private randomization that j is meant to conduct in part 6c, recall that the situation is that

i is in state P, j is in state G, and the favor is of some type $(x, y) \in S_{ij}$. Let \mathcal{I} denote the set of players other than i who are currently in state P. Then j 's private randomizing probabilities are $x/(v_i^G(\mathcal{I}))$ for " i is restored to grace," and the complementary probability for " i is not restored to grace." If our assertion that $v_i^G(\mathcal{I}) \geq \frac{(\mathcal{A}_i(S) - \mathcal{B}_i(S))}{r}$ is correct, then since $x \leq M_{ij}(S) \leq M_i(S)$, Condition (a) tells us that this probability is between 0 and 1.

Will i do the favor, if she is in state P, j is in G, and the favor type $(x, y) \in S_{ij}$? If she doesn't do it, she remains in P, with continuation value $v_i^P(\mathcal{I})$. If she does, her expected payoff is

$$-x + \frac{x}{v_i^G(\mathcal{I})} \cdot v_i^G(\mathcal{I}) + \left[1 - \frac{x}{v_i^G(\mathcal{I})} \right] \cdot \left[\psi_j(\mathcal{I}, i) v_i^P(\mathcal{I}) + (1 - \psi_j(\mathcal{I}, i)) \cdot v_i^P(\mathcal{I} \cup \{j\}) \right],$$

where $\psi_j(\mathcal{I}, i)$ is the public-randomization probability of j being sent to purgatory in these circumstances that is still to be specified for part 6c. To explain, in the display we have the cost of the favor, plus the probability that she is restored by j to G, times her continuation value there, times the probability that j does not restore her to G, times the quantity: her expected continuation value, depending on whether the public randomization causes j to join her and the members of \mathcal{I} in state P.

Suppose $v_i^P(\mathcal{I}) = v_i^P(\mathcal{I} \cup \{j\}) = 0$. Then the displayed expression is 0, which is $v_i^P(\mathcal{I})$, and i is indifferent between doing the favor or not. We assert that with j randomizing as specified, $v_i^P(\mathcal{I})$ is indeed 0, for every i and for every \mathcal{I} . That is, this form of purgatory is constructed so that its continuation value is 0. This is so because, once in P, i is getting no favors. The only way she can escape is by doing favors for players in state G that are mandated by S . But if she does one of these favors for, say, j , j randomizes between letting her back into G and keeping her in P, with a probability that is set so that her expected benefit from getting back to G just covers the cost of the favor she is called on to do. In equilibrium, she does the favor; but she is indifferent between doing so or not. Because she is always indifferent (between doing favors called for under S), her expected payoff is the same if she chooses not to do these favors; then she never does

a favor, never escapes P, never gets a favor, and consequently has expected payoff of zero.¹³

As for doing favors while in P that are not mandated by S , part 6b says that j , the favor receiver, will definitely make the public declaration that keeps her in P; she receives no compensation for doing such a favor and so, of course, she doesn't do it.

Turning to favors an i is called on to do while in state G (for others who are in state G), since the expected payoff of being in purgatory is zero, as is the expected payoff of autarky, Condition (a) implies that a player in state G will certainly do all favors called for by S (for others who are also in G), as the alternative is a lottery between being sent to P or moving to autarky forever.

What about the public randomizations in parts 4, 6b, and 6c? In part 4 and 6c, we need j to be indifferent between his two possible cheap-talk declarations, since the strategies call for him to use them both. (In part 4, circumstances dictate which declaration he uses; in part 6c, he employs a mixed [behavioral] strategy.) We have asserted that he weakly benefits by sending i to purgatory in part 4 and by keeping i in purgatory in part 6c. Therefore, to make him indifferent, if he sends i to purgatory in part 4 or keeps i in purgatory in part 6c, we mix with an outcome that is bad for him. In part 4, that outcome is autarky for all; we can't use this device in part 6c for perfection considerations, but we can send j to state P, which is just as good (that is, just as bad, for j).¹⁴ This then allows us to compute: in part 4, if \mathcal{I} is the set of folks in P when i fails to do a favor for j that she was meant to do, and j renounces her, the probability that we do *not* move to autarky is $\frac{v_j^G(\mathcal{I})}{v_j^G(\mathcal{I} \cup \{i\})}$. And in part 6c, where i has done a favor mandated under S for j but j has declared that i is not restored to grace, the probability that j is not sent to P is given by the same $\frac{v_j^G(\mathcal{I})}{v_j^G(\mathcal{I} \cup \{i\})}$, where now \mathcal{I} is the set of people

13. We are employing here—and elsewhere in the proof—the general result from the theory of dynamic programming that, for well-behaved problems (such as here, with discounting and bounded rewards), if a player is indifferent between following strategy x for a while and then reverting to strategy y , or just following strategy y , then strategies x and y give the same expected payoff. See, Kreps (2013, Proposition A6.10a and b).

14. We could equally well, in part 4, move j to P instead of moving everyone to autarky; this has the same impact on j 's incentives to issue his two possible cheap-talk declarations.

other than i who are currently in P . Under our assertion that $0 < \frac{A_i(S) - B_i(S)}{r} \leq v_j^G(I) \leq v_j^G(I \cup \{i\})$, these are indeed probabilities.

This leaves the public randomization in part 6b. The circumstances here are that i is in P , j is in G , and i did a favor for j that is not in S_{ij} . The declaration issued by j is sure to keep i in P , but because players other than i and j can't distinguish between situation 6b and 6c (they are distinguished by whether $(x, y) \in S_{ij}$ or not, which only i and j know), what happens to j under a public randomization if he says " i is not restored to grace" must be the same in the two cases.

What about our assertion that $v_j^G(I) \geq v_j^G(J)$ for $J \subset I$? Suppose that all players except for j follow the strategies outlined above, while j acts as above *except that* whenever he is called on (in step 6) to announce whether some i is released from purgatory, he always releases i . The public randomization that ensues is constructed so that he is indifferent between the specified strategy (where he randomizes between the two messages) and this alternative strategy (where he always releases i), so following the alternative strategy provides him the same expected payoff. Then, under this alternative strategy for him (with everyone else following the specified strategies), there is no chance that he will enter P , if he begins in G . He will, from any starting point where he is in G , receive all favors and only those favors mandated by S . But if a rival is in P , he benefits (weakly) by not having to do favors for that rival, at least for a while. Note, if i is added to the list of (other) people in P , then i 's presence in P can act to draw other i' into P as well (as they follow the mandates of part 6c of the strategies), which increases the (weak) benefit to j . This establishes our assertion.

This implies that we do, indeed, have a perfect social equilibrium. Except that this argument is not quite a proof, as it contains a circularity that must be addressed. To identify the probabilities in parts 6b and 6c (both the randomizing probability employed by the favor receiver and the public-randomization probability used if a favor is given but the favor receiver does not restore the favor giver to grace), we employed the various value functions for players who are in state G . But those value functions depend on the probabilities employed.¹⁵ This is a fixed-point problem: we need probabilities that generate values that, in

15. This is not true of the public-randomization probabilities in part 4, because that randomization is never reached along the path of play from *any* initial position.

turn, give back the probabilities. The argument that such probabilities and corresponding value functions exist is provided in the Appendix. ■

VI. COMPARING EX ANTE AND EX POST CHEAP TALK

Propositions 1 and 2 tell us that at least insofar as implementing time-homogeneous selections are concerned, everything that can be done with ex ante cheap talk for a given r can be done with ex post cheap talk. More precisely, every selection S that passes an obvious incentive-compatibility test, Condition (a), can be implemented as a perfect social equilibrium with ex post cheap talk; since Condition (b) is required for ex ante cheap talk to work, ex post cheap talk seems superior.

We urge some caution in drawing this conclusion, on two grounds:

- (i) The equilibrium strategies in Proposition 1 are quite simple and straightforward; the off-path strategies in Proposition 2 make sense if you stare at them long enough, but they cannot be called simple and straightforward.
- (ii) As long as Conditions (a) and (b) in Proposition 1 hold with strict inequalities, the equilibrium in Proposition 1 can be constructed with strict best responses for every player in every situation. It is part and parcel of the equilibrium in Proposition 2 that the favor receiver is indifferent whether to forgive a favor that is not done along the path of play, that (off the path of play) favor receivers use randomized strategies, that they are indifferent between allowing the favor giver back into a state of grace or not, and that a favor giver who is in state P is indifferent whether to do the favors called for under S . (Suppose a favor giver in state G fails to do a mandated favor and then quietly warns the favor receiver, "If you denounce me and I go to P, I will not do any favors for you." If credible, this threat causes the equilibrium to collapse; it is essential that she does those favors, so that favor receivers is [weakly] better off with her in state P.) To the extent that all these indifferences make the equilibrium delicate, the equilibrium in Proposition 1 is better, when it works.

In this regard, it is possible, with *ex post* cheap talk, to construct equilibria that implement some selections S without all the indifferences required in the equilibrium of Proposition 2. The first sort of indifference—a favor receiver who does not receive a favor along the path of play is indifferent in what he says—cannot be avoided; the argument given at the start of Section V shows this. But by constructing other forms of punishment, one can create out-of-equilibrium strategies that have much more robust (that is to say, strict) best responses. These equilibria, however, require that there is a lot of “slack value,” in the sense that the left side of the inequality in Condition (a) is substantially less than the right-hand side; moreover, they require that inequalities of this sort hold where in place of the left-hand side we have the most expensive favors that a player might ever do, in the all-favors selection. In the Online Appendix, we provide details of some of these equilibria.

Of course, the ultimate test for which of *ex ante* or *ex post* cheap talk is employed is empirical. It is clear to us that both are employed in different contexts. The most we can say here is that perhaps this analysis suggests how the two might work, which can then be used to guide empirical investigation.

VII. EFFICIENCY AND MIXTURES OF SELECTIONS

The space of feasible utility imputations (vectors of *ex ante* expected payoffs for each player) is clearly a convex set in R^I , so standard methods (see, for instance, Kreps 2013, Proposition 8.10) imply that every (first-best) efficient outcome of the game can be derived as the solution of a maximization problem, where the objective function is a weighted sum of individual expected payoffs, for nonnegative weights. Conversely, every such solution is efficient, as long as the weights are strictly positive. For the remainder of this discussion, we will restrict attention to weighting vectors α that are normalized to sum to one; that is, that lie in the $I - 1$ -dimensional unit simplex.

Letting $\alpha = (\alpha_i)_{i=1, \dots, I}$ denote a vector of nonnegative weights, it is clear that an outcome of the game maximizes the α -weighted sum of expected utilities of the players if and only if (a) all *i*-for-*j*

favors with cost–benefit pair (x, y) such that $\alpha_j y > \alpha_i x$ are done, and (b) no favors are done if the reverse strict inequality holds. Favors for which $\alpha_j y = \alpha_i x$ can be done or not, with no impact on the objective function.

Since we have assumed that the supports of the distributions of cost–benefit pairs for all i and j are finite, there are finitely many possible selections. Hence we know that the space of feasible utility imputations is a polytope, the convex hull of the utility imputations generated by the finitely many selections. Therefore the efficient frontier is the northeast boundary of this polytope, with vertices corresponding to selections for which some nonnegative weighting vector α gives the selection according to the rule that for the ordered pair i, j , an i -for- j favor with cost–benefit vector (x, y) is necessarily in the selection if $\alpha_j y > \alpha_i x$ holds and is not in the selection if the reverse strict inequality holds. Faces of the efficient frontier are given by weighting vectors where $\alpha_j y = \alpha_i x$ for at least some i, j pairs and some i -for- j favors (x, y) , and every point on a face of the polytope—every efficient outcome, therefore—can be obtained as a simple convex combination of the outcomes obtained from pure selections.

For “most” weighting vectors α (that is, for a set of full measure in $I - 1$ -dimensional unit simplex), there are no “ties,” hence implementing the corresponding efficient point comes down to the question of whether the corresponding (unique) selection can be implemented; our results answer this question directly. But “most” points on the efficient frontier correspond to those weighting vectors for which there are ties, and the implementation of those points requires more thought.

Consider the following more general question. Let \mathcal{M} denote the space of all probability distributions on Σ , the space of all selections, and let μ denote a generic element of \mathcal{M} , with $\mu(S)$ denoting the probability of selection S under μ . Finally, let $\text{supp}(\mu)$ be the support of μ . Then for given $r > 0$, when can a perfect social equilibrium be constructed (involving either ex ante or ex post cheap talk) that provides for player i the payoff

$$v_i(\mu, r) := \sum_{S \in \text{supp}(\mu)} \mu(S) \frac{\mathcal{A}_i(S) - \mathcal{B}_i(S)}{r} ?$$

There are two more or less obvious ways to try to implement μ in this sense. If each $S \in \text{supp}(\mu)$ satisfies Condition (a), then we can implement μ with ex post cheap talk by conducting an initial

public randomization that chooses some selection S according to μ , and then implementing the chosen S via the methods of Proposition 2. And if each $S \in \text{supp}(\mu)$ satisfies Condition (b) as well, we can use an initial randomization and the methods of Proposition 1 to implement μ with ex ante cheap talk.

But if a particular convex combination of vertices involves selections that cannot be implemented, this does not necessarily imply that the convex combination cannot be implemented. A caricature of our problem, with only two players and degenerate distributions for cost–benefit pairs, makes the point. Suppose that 1-for-2 and 2-for-1 favors arrive at the same rate: $\lambda_{12} = \lambda_{21} = 1$. Suppose that 1-for-2 favors always have a cost of 4 and provide a benefit of 8, and 2-for-1 favors have cost 3 and benefit 3. There are three “pure,” non-null selections, namely, only the 1-for-2 favors, only the 2-for-1 favors, and both. Obviously, the selections that involve only one player doing favors for the other cannot be implemented. If favors arrive at equal rate, the expected flow-rate of costs to player 1 exceeds the expected flow-rate of benefits, if all favors are done, so this selection cannot be implemented for any interest rate. But suppose we try to implement a situation in which 1-for-2 favors are done half the time, and 2-for-1 favors are always done. This gives a positive net flow-rate of benefits for both players, hence for r close enough to zero, the expected value of an equilibrium with this outcome provides enough value to each player to cover the cost of any immediate favor. Because of the additive structure of rewards, we can implement this “mixture” of selections ($\mu(S^A) = 0.5$ and $\mu(S^1) = 0.5$ where S^1 is the selection in which only 2-for-1 favors are done): each time a favor opportunity arises, and before anything else happens (concerning that favor opportunity), a public randomization is conducted that determines whether we follow the rules for implementing S^1 or S^A . The point is, by re-randomizing at each favor opportunity, the continuation values of the players, at least along the equilibrium path, after the current encounter is over, remain $v_i(\mu, r)$.

Obviously, this strategy of re-randomizing at each favor opportunity generalizes. As long as

$$\max_{S \in \text{supp}(\mu)} M_i(S) \leq v_i(\mu, r)$$

for each i , we can implement μ with ex post cheap talk, by employing the equilibrium of Proposition 2, supplemented with this

procedure of re-randomizing over the selections according to μ at every arrival of a favor opportunity. We can do so as well with ex ante cheap talk, if (in addition) the obvious analogue to Condition (b) holds for μ : for all i and j , if (x, y) is an i -for- j favor that is not done in any one of the selections in the support of μ , then $x \geq \max_{S \in \text{supp}(\mu)} \mathcal{M}_{ij}(S)$.

Neither of these two ways of implementing mixtures of selections dominates the other. The example just given shows that the re-randomizing procedure may work when the randomize-at-the-start procedure may fail. Conversely, if one is aiming for ex ante implementation, two selections S and S' may both satisfy Condition (b) individually, but for any nontrivial mixture of them, the analog to Condition (b) may fail.

VIII. EXTENSIONS AND VARIATIONS

Although Propositions 1 and 2 paint a nice and concise theoretical picture, several extensions and variations are available that fill out the theoretical picture and that allow us to begin to move from our stylized model to models that capture more of the institutional features of real-life models.

VIII.A. Folk Theorems and Anti-Folk Theorems

Our attention here has been on the question, which selections can be implemented for given $r > 0$? It is somewhat traditional in literature of this sort to ask instead, which selections can be implemented, or even *approximately implemented*, for “small enough” r ? The italicized phrase *approximately implemented* needs a definition: we want to know if there are perfect social equilibria for small r whose normalized payoffs converge to the normalized payoffs from S , as r approaches zero. That is, if $\{r^n\}$ is a sequence of interest rates with limit 0, can we find perfect social equilibria, one for each r^n , with respective expected payoffs given by v_i^n for player i , such that $r^n v_i^n \rightarrow A_i(S) - B_i(S)$?

Fix a selection S . If $A_i(S) < B_i(S)$ for any player i , then it is clear that we cannot implement S for any r , nor can we approximately implement S in the sense just described. Ignoring knife-edge cases for which $A_i(S) \geq B_i(S)$, with some equalities, we have the following obvious corollary to Propositions 1 and 2.

COROLLARY 1. Suppose that for a given selection S , $A_i(S) > B_i(S)$.

Then for some $r(S) > 0$, S can be implemented in a perfect social equilibrium with ex post cheap talk for all $r \leq r(S)$. If, in addition, S satisfies Condition (b) (which, it may be worth noting here, is independent of the value of r), the same conclusion is true for ex ante cheap talk.

Still ignoring the knife-edge cases, this leaves open one question: suppose $A_i(S) > B_i(S)$ for all i , but Condition (b) fails. Can selection S be *approximately implemented*, for small interest rates, with ex ante cheap talk? We conjecture that the answer is yes, if no restrictions are put on the richness of the cheap-talk language that can be employed; we explain why we believe this conjecture is true in the Appendix. But *if* the cheap-talk language is limited—and if we are looking for “simple” implementations of outcomes, which could mean (among other things) a limit to the vocabulary of cheap talk—we get the following “anti-folk theorem”:

PROPOSITION 3. Suppose we restrict potential favor receivers to cheap talk only before the favor decision is made and, moreover, the favor receiver is limited to his choice of two (cheap-talk) messages before i acts, such as that the favor giver is excused or not. Suppose that for some pair i and j with $\lambda_{ij} > 0$, there are four cost-benefit pairs, $(x(k), y(k))$ for some i and j with

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

Then the u-efficient selection S^U cannot be “implemented asymptotically”: there exists a strictly positive constant K such that, if (v_1, \dots, v_I) is a vector of expected payoffs to the players in any perfect social equilibrium for interest rate r ,

$$r \sum_i v_i \leq \sum_i [A_i(S^U) - B_i(S^U)] - K.$$

That is, the sum of normalized equilibrium payoffs in any perfect social equilibrium is uniformly bounded away from the sum of normalized payoffs from the selection S^U .

The proof is provided in the Appendix.

VIII.B. *Time Homogeneity and Unilateral Discretion*

In Section VII and the preceding subsection, we have begun to move away from the time-homogeneous implementation of a single selection S . In discussing mixtures of (pure) selections, we don't move too far away: first we propose an initial public randomization followed by the implementation of a pure selection; then we consider strategies in which there is re-randomization among the selections at each favor opportunity, but in time-homogeneous fashion, so (at least) continuation values after every favor opportunity (along the path of play) never change. But when we think about the approximate implementation of a selection for small r , we are contemplating—in theory, at least—strategies in which play after a favor opportunity depends on what has happened (and is common knowledge) at that favor opportunity.

In this regard, recall the brief mention of unilateral discretion back in Section III. Even with no cheap talk by favor receivers, perfect social equilibria can be constructed in which the favor giver does not do all favors but exercises her own discretion in which favors she chooses to do. In these equilibria, it is essential that the continuation payoff for the favor giver after she does a favor exceeds her continuation payoff if she doesn't do a favor; were they the same, she would do no favors. This implies two things. First, what happens along the path of play cannot be time homogeneous. Second, the favors she will choose to do will be determined entirely by their cost (that is, in any social equilibrium), since her continuation values can only depend on whether she does the favor. Hence if we look for such unilateral-discretion equilibria, considerations in the spirit of Condition (b) from Proposition 1 intrude. Unilateral discretion can only involve selections that satisfy Condition (b). But if one has a selection that satisfies Conditions (a) and (b) with strict inequalities, a folk theorem without cheap talk can be shown: as the interest rate r approaches 0, equilibria can be constructed that asymptotically give the expected payoff of the selection. Conversely, if Condition (b) fails for the selection S^U , an anti-folk theorem along the lines of Proposition 3 can be proved. For details, see the Online Appendix.

VIII.C. *Varying the Informational Endowments*

We have supposed throughout that i and j both know the precise values of cost and benefit of all i -for- j favors and that all

the other players know when favor opportunities arise and whether the favor has been done. The folk theorem just alluded to can be read as saying: if a selection S satisfies Conditions (a) and (b) with strict inequalities, and if the players are quite patient, then it is (asymptotically) unnecessary that the favor receiver of any favor know anything more than does the general population of players, and the favor giver only needs to know the value of x , her cost of doing the favor. Two other “informational variations” on our propositions constitute low-hanging fruit:

- (i) Suppose that i and j both know (x, y) for every i -for- j favor opportunity, but other players are generally unaware that a favor opportunity has occurred unless the favor is done. The equilibria in Propositions 1 and 2 can still be made to work. In the equilibrium of Proposition 1, if an i -for- j favor opportunity arises for which $(x, y) \in S_{ij}$, j can broadcast, “Player i has the opportunity to do me a favor, which I expect her to do.” It is necessary that other players can then see if the favor is done, but as long as they can see this, the equilibrium goes through. The key to this is that j is punished for a favor that i could do for her but does not do (and for which i is not excused); the same punishment would stop j from making the broadcast just suggested, if the favor is one that i will not do.

For the equilibrium of Proposition 2: if an i -for- j favor opportunity arises, which i does not do, if i and j are in G , and the favor should have been done (is in S), then j can broadcast, “Player i had the opportunity to do a favor for me, should have done so, and didn’t do it.” Then we have a public randomization in which the two outcomes are that i alone goes to state P , or both i and j go to P ; since this randomization makes j indifferent between talking and remaining silent, we can suppose that he speaks. Things are a bit trickier when i is in P : if an i -for- j favor opportunity arises which i doesn’t do, j is perfectly happy (ex post) to let it ride; but if i does a favor, the equilibrium requires that j make some declaration, although he would prefer not to, so we require that other players observe that i did the favor and it is j ’s turn to speak.

- (ii) In the equilibrium of Proposition 1, it is not necessary that both i and j know the cost and benefit levels of

each i -for- j favor; they only need to know whether the favor type (x, y) is or is not in S . For the equilibrium of Proposition 2, things aren't quite so simple. Off the path of play, to achieve the correct randomization, a favor receiver must know the exact value of x . But favor givers can get by with only knowing if a given favor is or is not part of the selection S .

These are easy extensions; more interesting questions, which require serious analysis, concern how robust our constructions are to other variations in information endowments. A few simple results of this sort are given in the Online Appendix.

VIII.D. *Endogenous Informational Endowments*

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 might want to invest in and then employ local information. This "investment in information" is, of course, something we see in the real world: countries invest in embassies and diplomacy, trade missions, and multinational 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.

VIII.E. *Externalities*

We've looked here at the case where, when an i -for- j favor opportunity arises, only i and j are potentially affected directly by whether the favor is done. Our focus on this case is based on the premise that when information is local, it is most likely to be held by the two parties most directly concerned with whether the favor is done. But in at least some of the applications we have in mind (e.g., organizational behavior in work settings), favors can generate externalities for the other parties in the "game."

To say a bit more here about the application to work settings, imagine that we have an organization consisting of a collection of individuals. Instead of thinking of the random arrival of i -for- j

favors, think of the random arrival of tasks for the various players i . Each task, if done, has a cost x for the player i who does the task, and it generates a vector of positive benefits for all the other players. Imagine that some distinguished player, say player 1, who might be called “boss” or “department chair” or “dean” or “manager,” is able to tell, when a task-for- i arises, what is x and what is the vector of benefits for all the other players, including himself. (In addition, i must have this information.) Then, if we reinterpret a task-for- i as an i -for-1 favor with externalities, we are back in the setting of our model (albeit with externalities). Note well, we do *not* require that the “referee” for which tasks should be done or should have been done—that is, player 1—is neutral to whether the task is done. Player 1 can receive external benefits if the task is done.

So it is interesting to ask: how do externalities affect our equilibrium constructions?

If we suppose that both i and j know which i -for- j favors ought to be done and which not (which, if favors generate externalities, could involve more than i 's costs and j 's benefits), then the equilibrium of Proposition 1 works nearly without change. The only modification is that in assessing whether Condition (a) holds for each player i , one needs to compute the equilibrium net benefits that i gets in equilibrium, including net benefits from externalities from favors in which i is otherwise unconcerned.¹⁶

Things are not so pleasant with the equilibrium of Proposition 3 (hence this is another reason not to dismiss *ex ante* cheap talk) To give a sense of what can go wrong, suppose an i -for- j favor opportunity arises which i is meant to do according to S , but which i doesn't do. In the equilibrium, j should denounce i , and j is willing to do so, because having i in state P is weakly better for j . But suppose that some third-party k does favors for i that generate particularly nice positive externalities for j . If j sends i to P, k stops doing favors for i . So it may be worse for j to have i in P than in G, which (of course) causes the equilibrium construction to collapse.

16. Note in this regard the special case where player 1 is always treated as the favor receiver, because he is a specialist in assessing the costs and vector of benefits for each task opportunity. If player 1 announces, “This is a task that i ought to do,” and then i fails to do it, both i and player 1 should be punished. Readers can judge from personal experience if department chairs and deans are dealt with in this fashion. Of course, if player 1 has tasks to do other than keeping tabs on others, someone must monitor her.

One of the alternative ex post cheap-talk equilibria we construct in the Online Appendix doesn't fall afoul of these problems, at least in the case where all externalities are positive. See the Online Appendix for details.

VIII.F. *Transferable Utility and Corruption*

We began by ruling out transfer payments between the parties; because utility is transferable to some extent in the real world, we conclude with a few remarks on this point. If utility is transferable among the parties on a one-to-one basis, and absent the presence of externalities, the point of this article is lost. When an *i*-for-*j* favor opportunity arises, where both *i* and *j* know the cost to *i* and the benefit to *j*, one expects them to come to some mutually agreeable, bilateral deal. Assuming the cost to *i* is less than the benefit to *j*, *j* can offer to pay *i* her cost; or *i* can demand from *j* a payment equal to his benefit. If the cost to *i* exceeds the benefit to *j*, there is no reason for the favor to be done.¹⁷ The issue becomes one of an immediate bargain being struck.

But in some instances, monetary transfers are not possible and, even where they are possible, they may be inefficient; that is, it may cost an *i* more in utility than she can deliver to *j*, in a direct transfer. Our assumption of no transfers is, of course, extreme. But it is not entirely unreasonable.

Adding together the possibility of transfer payments and externalities gives rich ground for further analysis, especially if the transfer payments can be made surreptitiously. Consider, in this regard, the equilibrium of Proposition 1. An *i*-for-*j* favor opportunity arises, that will cost *i* the amount x and provide benefits $y < x$ to *j*, but that generates further (positive) externalities for all the other players in amounts large enough so that on efficiency grounds, this is a favor that ought to be done. If *i* can approach *j* before *j* broadcasts his ex ante message and offers him, say, $\frac{x+y}{2}$ if he will excuse this favor, *j* is better off accepting than refusing to issue the excuse. So, in particular, if we enlist the specialization

17. Think of the circle and a case where all *i*-for-*i* + 1, for $i = 1, 2, \dots, I - 1$, have cost 1 and benefit 2, while the *I*-for-1 favors have cost 2.1 and benefit 2. You might at first think that we need for this last, u-inefficient favor to be done, to keep 1 happy and willing to do favors for 2, and so on. But with transferable utility, this simply isn't so. Player 2 pays player 1 for 1-for-2 favors, which is all the compensation 1 needs.

in which some distinguished player, the department chair, say, is always in the position of favor receiver/referee, we begin to see a case (insofar as surreptitious transfers are possible) for constructing and enforcing sanctions against corrupt practices. (Or, we see why some decisions of this sort might be moved from the level of, say, a department chair—who may be more susceptible to under-the-table side payments of various sorts—to a decanal level.)

APPENDIX

A. The Fixed-Point Argument for the Proof of Proposition 2

In the proof of Proposition 2 given in the text and, in particular, in the initial description of strategies for the alleged equilibrium, the specification of three types of probabilities is left to be determined endogenously:

- *In part 4 of the description of the equilibrium strategies, the probability that play moves to autarky, if i in G fails to do a favor for j and j denounces i , depending on the list of other folks who are in purgatory:* If \mathcal{I} denotes the folks already in purgatory, this probability $\pi_j(\mathcal{I}, i)$ should satisfy

$$\pi_j(\mathcal{I}, i) \cdot 0 + (1 - \pi_j(\mathcal{I}, i)) \cdot v_j^G(\mathcal{I} \cup \{i\}) = v_j^G(\mathcal{I}),$$

$$\text{hence } \pi_j(\mathcal{I}, i) = 1 - \frac{v_j^G(\mathcal{I})}{v_j^G(\mathcal{I} \cup \{i\})},$$

so that j is indifferent between sending i to purgatory and thereby triggering this public randomization, or not doing so.

- *In part 6c, the probability that j (in state G) does not release i (in P) from purgatory if i does an appropriate favor for j when i is in purgatory and j is in grace:* If the cost of the particular favor is x and the folks other than i who are in purgatory is \mathcal{I} , this probability $\rho_i(x, \mathcal{I})$ must satisfy

$$x = (1 - \rho_i(x, \mathcal{I}))v_i^G(\mathcal{I}), \text{ hence } \rho_i(x, \mathcal{I}) = 1 - \frac{x}{v_i^G(\mathcal{I})};$$

that is, it must make i just willing to do the favor.

- In parts 6b and 6c, the probability that j is sent to purgatory (by public randomization) if j does not release i from purgatory after i does a favor for j , when i is in purgatory and j is not, depending on the folks \mathcal{I} besides i who are in purgatory: Calling this probability $\psi_j(\mathcal{I}, i)$, it must satisfy

$$(1 - \psi_j(\mathcal{I}, i))v_j^G(\mathcal{I} \cup \{i\}) = v_j^G(\mathcal{I}), \text{ hence } \psi_j(\mathcal{I}, i) = 1 - \frac{v_j^G(\mathcal{I})}{v_j^G(\mathcal{I} \cup \{i\})}.$$

The ambiguity here is that these probabilities depend on the value functions v^G , which in turn depend on these probabilities. Being more precise about this, the continuation-value functions (from various starting points in the game) depend on the ρ and ψ probabilities; the π probabilities are irrelevant to the continuation values, because the π (part 4) randomization never occurs along the path of play from any point in the game tree which matters to the computation of a continuation-value function, if players follow the prescribed strategies. But if we look at points in the game tree where some players are in purgatory, the ρ and ψ probabilities are used along the path of subsequent play and do affect the continuation values for various points at which players begin in purgatory.

Hence the ρ and ψ probabilities, together with the value functions, represent a fixed point. To prove the existence of probabilities and values that fit together (which then completes the proof of the proposition), reason as follows.

Write \mathcal{P} for the domain of these probabilities. (That is, a point in \mathcal{P} specifies values for the ρ and ψ probabilities.) We define a function $\Phi : \mathcal{P} \rightarrow \mathcal{P}$ in two steps.

First, given a specification of the ρ and ψ probabilities (that is, an element of \mathcal{P}), for fixed i , let $w_i(\mathcal{I})$ denote the expected value to player i of being in state G , if the roster of players in P is given by \mathcal{I} (so this is for $i \notin \mathcal{I}$) and, with the exception given below, play is according to the equilibrium described in the proposition, using for ρ and ψ probabilities the values initially specified. The one exception is that, if any player j who is in purgatory does a favor for i , i simply releases j from purgatory (whether the favor was warranted by the selection or not).

The impact of this one change is that when computing $w_i(\mathcal{I})$, there is no chance that i (who is in G) will be dragged down into P ;

since i releases all j from P , no public randomization is conducted that could result in i going to P .

We assert three things about these “pseudo-value functions” $w_i(\mathcal{I})$:

- (i) $\frac{A_i(S) - B_i(S)}{r} \leq w_i(\mathcal{I}) \leq \frac{A_i(S)}{r}$;
- (ii) if $\mathcal{I} \subset \mathcal{J}$ (and $i \notin \mathcal{J}$), then $w_i(\mathcal{J}) \geq w_i(\mathcal{I})$; and
- (iii) the pseudo-value functions $w_i(\mathcal{I})$ change continuously in the initially specified ρ and ψ probabilities.

The first and third of these assertions are straightforward. The first follows because, along the path of pseudo-play (which now gives no chance that i would be dragged down into purgatory), the flow rate of expected benefits to i always lies between $A_i(S)$ (if i only gets favors mandated by S) and $A_i(S) - B_i(S)$ (if i has to do all mandated favors; that is, if no one is in purgatory). The third simply involves noting that over any finite horizon, a small change in the initial specification of the ρ and ψ probabilities gives a small change in the evolution of the system, and the distant horizon is irrelevant, because of discounting and a bound on the expected flow-rate of payoffs.

But assertion 2 needs a bit of explanation. The explanation is essentially given in the text, but to reiterate: it says that the more (other) people begin in purgatory, the (weakly) better off is i . Having others in purgatory is clearly better for i in the short run: the flow of favors to i doesn't change, while the flow of favors i must do is reduced, the greater the set of folks in purgatory. And the more folks there are in purgatory, the greater is the chance that some third party will be dragged into purgatory (by the working out of part 6c of the specification of strategies). Finally, because of the way in which we modified i 's strategy when computing her pseudo-value function, there is no chance that she will be dragged into purgatory. (If we hadn't made that modification, having more people there increases the odds that she will be dragged down via part 6c, which could be worse for her overall.) This gives us 2.

Having computed these pseudo-value functions, compute the new values of the ρ and ψ probabilities, as follows:

$$\rho_i(x, \mathcal{I}) = 1 - \frac{x}{w_i(\mathcal{I})} \quad \text{and} \quad \psi_j(\mathcal{I}, i) = 1 - \frac{w_j(\mathcal{I})}{w_j(\mathcal{I} \cup \{i\})}.$$

Note that these are indeed probabilities by virtue of 1 and 2.

We assert that Φ is a continuous map from \mathcal{P} into \mathcal{P} ; assertion 3 tells us that the value functions are continuous in the originally specified probabilities, and 1 and simple math imply that the new probabilities are continuous in the value functions. Therefore, by Brouwer's fixed-point theorem, ϕ has a fixed point.

We assert, finally, that this fixed point is indeed an equilibrium to our game, even when we go back to the equilibrium specification for what each j does if he is in G and a favor (warranted by \mathcal{S} or not) is done for him by, say, i who is in P : now, instead of releasing i from purgatory, he steadfastly refuses to do so if the favor is not prescribed by \mathcal{S} , and he randomizes about what to do (using $\rho_i(x, \mathcal{I})$) if the favor is in the selection \mathcal{S} . Hence there is a chance that he will be subjected to a public randomization that could result in him being dragged into purgatory.

But this public lottery is constructed so that j is indifferent at every point where he must make a release decision between allowing i to escape purgatory and refusing to release her. Applying the general argument of footnote 13, j 's expected value in G of using his actual strategy is the same as if he uses the "pseudo-strategy" (as long as the value function is fixed), which completes the proof. ■

B. The Proof of Proposition 3

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

$$\sum_i \left[\sum_{j \neq i} \lambda_{ji} A_{ji} - \lambda_{ij} B_{ij} \right] = r \sum_i \sum_j \sum_{n=1}^{\infty} \mathbf{E} \left[e^{-rT_{ij}^n} (y_{ij}^n - x_{ij}^n) 1_{\{x_{ij}^n < y_{ij}^n\}} \right],$$

where T_{ij}^n is the random arrival time of the n th i -for- j favor and (x_{ij}^n, y_{ij}^n) is the cost-benefit vector for this favor. (If $\lambda_{ij} = 0$, either omit j from the second sum for this i or take $T_{ij}^n \equiv \infty$.) On the other hand, the normalized sum of payoffs in the equilibrium is

$$r \sum_i \sum_{j \neq i} \sum_{n=1}^{\infty} \mathbf{E} \left[e^{-rT_{ij}^n} (y_{ij}^n - x_{ij}^n) 1_{\{i \text{ does this favor for } j\}} \right],$$

where the expectation being taken involves the exogenously determined random arrival times of favors, the exogenously determined random 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

$$r \sum_i \sum_{j \neq i} \sum_{n=1}^{\infty} \mathbf{E} \left[e^{-rT_{ij}^n} (y_{ij}^n - x_{ij}^n) [1_{\{x_{ij}^n < y_{ij}^n\}} - 1_{\{i \text{ does this favor for } j\}}] \right]. \quad (3)$$

(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 $y_{ij}^n > x_{ij}^n$, the first indicator function is 1, so the term inside the smaller square brackets is either 1 or 0; if $y_{ij}^n < x_{ij}^n$, 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.

The proposition posits that for some ordered pair i and j , there are four cost-benefit vectors $(x(k), y(k))$ (in the support of the distribution of cost-benefit vectors for i -for- j favors) such that $y(1) > x(1) > x(2) > y(2) > y(3) > x(3) > x(4) > y(4)$. Refer to an i -for- j 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 -for- j favor, the cost-benefit vector makes this a favor of type k .

Now go back to the difference in (3), and examine the term for this specific i and j (ordered) pair. That is, we are looking at

$$(4) \quad r \sum_{n=1}^{\infty} \mathbf{E} \left[e^{-rT_{ij}^n} (y_{ij}^n - x_{ij}^n) [1_{\{x_{ij}^n < y_{ij}^n\}} - 1_{\{i \text{ does this favor for } j\}}] \right].$$

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, 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

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

Proposition 3 is established, for $K = \lambda_{ij} K^1$.

The proposition posits that j can only issue one of two cheap-talk messages, which we denote by A and B . We use v_{AF} for i 's continuation value if j issues message A and i does the favor, v_{AN} for i 's continuation value if j says A and i does not do the favor, and similarly for v_{BF} and v_{BN} . 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.

- (i) 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.)
- (ii) But if i does the favor with positive probability when the favor is of type 1, we know that $v_{AF} - x(1) \geq v_{AN}$. Since $x(1) > x(2) > x(3) > x(4)$, this implies that $v_{AF} - x(i) > v_{AN}$ for $i = 2, 3$, or 4 , which means that, if the favor is of type 2, 3, or 4, and j says A , i is certain to do the favor.
- (iii) 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 will 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.
- (iv) 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 $y(2) + u_{AF}$, whereas 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, completing the proof of Proposition 3. ■

It should be clear that the argument just given relies heavily on the assumption that the favor receiver is allowed to send only one of two messages. 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, as long as $\mathcal{A}_i(S^U) > \mathcal{B}_i(S^U)$ for all i . Indeed, we conjecture that if a selection \mathcal{S} satisfies $\mathcal{A}_i(\mathcal{S}) > \mathcal{B}_i(\mathcal{S})$ for all i , then with a sufficiently rich language of ex ante cheap talk by favor receivers, perfect social equilibria for small interest rates r can be constructed which, as $r \rightarrow 0$, have normalized payoffs that approach the normalized payoffs of \mathcal{S} . 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 transferable 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.

Suppose now that, instead of the u-efficient selection, we are looking at an "efficient" selection, in that sense that for some strictly positive weighting vector α , the selection is such that i -for- j favors are in selection if $\alpha_j y > \alpha_i x$ and are not in the selection if the reverse inequality holds. To avoid complications, assume that the weighting vector is such that there are never equalities in this comparison. Then the proof of the proposition is easily

modified to show that for any perfect social equilibrium, the *weighted* sum of normalized equilibrium payoffs is uniformly bounded away from the *weighted* sum of payoffs derived from this selection.

GRADUATE SCHOOL OF BUSINESS, STANFORD UNIVERSITY

SUPPLEMENTARY MATERIAL

An Online Appendix for this article can be found at QJE online (qje.oxfordjournals.org).

REFERENCES

- Barro, Robert J., "Reputation in a Model of Monetary Policy with Incomplete Information," *Journal of Monetary Economics*, 17 (1986), 3–20.
- Barro, Robert J., and David B. Gordon, "Rules, Discretion, and Reputation in a Model of Monetary Policy," *Journal of Monetary Economics*, 12 (1983), 101–121.
- Barth, Mary E., "Global Financial Reporting: Implications for U.S. Academics," *Accounting Review*, 83 (2008), 1159–1179.
- Bowen, Renee, "Forbearance in Optimal Multilateral Trade Agreements," Mimeo, Stanford Graduate School of Business, 2011.
- Bratton, William W., "Enron, Sarbanes-Oxley, and Accounting: Rules versus Principles versus Rents," *Villanova Law Review*, 48 (2003), 1023–1057.
- Crawford, Vincent P., and Sobel Joel, "Strategic Information Transmission," *Econometrica*, 50 (1982), 1431–1451.
- Ellison, Glenn, "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching," *Review of Economic Studies*, 61 (1994), 567–588.
- Freed, Daniel J., "Federal Sentencing in the Wake of Guidelines: Unacceptable Limits on the Discretion of Sentencers," *Yale Law School Faculty Schoolship Series*, Paper 2050, 1992; http://digitalcommons.law.yale.edu/fss_papers/2050.
- Fudenberg, Drew, Levine David, and Maskin Eric, "The Folk Theorem with Imperfect Public Information," *Econometrica*, 62 (1994), 997–1039.
- Ghosh, Parikshit, and Ray Debraj, "Cooperation in Community Interaction Without Information Flows," *Review of Economic Studies*, 63 (1996), 491–519.
- Hopenhayn, Hugo, and Hauser Christine, "Trading Favors: Optimal Exchange and Forgiveness," Meeting Papers 125, Society for Economic Dynamics, 2004.
- Kandori, M., "Social Norms and Community Enforcement," *Review of Economic Studies*, 59 (1992), 63–80.
- Kessler, Daniel P., "Evaluating the Medical Malpractice System and Options for Reform," *Journal of Economic Perspectives*, 25, no. 2 (2011), 93–110.
- Kreps, David M. *Microeconomics for Managers* (New York: Norton, 2004).
- , *Microeconomic Foundations I* (Princeton, NJ: Princeton University Press, 2013).
- Krishna, Vijay, and Morgan John, "The Art of Conversation: Eliciting Information from Experts through Multi-Stage Communication," *Journal of Economic Theory*, 117 (2004), 147–179.
- Kydland, F. E., and E. C. Prescott, "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, 85 (1977), 473–491.
- Lucas, Robert E., "Rules, Discretion, and the Role of the Economic Advisor," in *Rational Expectations and Economic Policy*, Stanley, Fischer, ed. (Chicago: University of Chicago Press, 1980), 199–210.

- Maggi, G., "The Role of Multilateral Institutions in International Trade Cooperation," *American Economic Review*, 89, no. 1 (1999), 190–214.
- Milgrom, Paul, and Roberts John *Economics, Organizations, and Management* (Englewood Cliffs, NJ: Prentice Hall, 1992).
- Mobius, Markus, "Trading Favors," Working paper, 2001.
- Modigliani, Franco, "The Monetarist Controversy, or Should We Forsake Stabilization Policy?," *American Economic Review* (March 1977), 1–19.
- Posner, Richard, "Rules versus Discretion—Posner's Comment," Becker–Posner Blog, September 16, 2007.
- Radner, Roy, Myerson Roger, and Maskin Eric, "An Example of a Repeated Partnership Game with Discounting and with Uniformly Inefficient Equilibria," *Review of Economic Studies*, 53 (1986), 59–69.
- Rosenthal, Robert W., "Sequences of Games with Varying Opponents," *Econometrica*, 47 (1979), 1353–1366.
- Schipper, Katherine, "Principles-Based Accounting Standards," *Accounting Horizons*, 17 (2003).
- Simons, Henry C., "Rules versus Authorities in Monetary Policy," *Journal of Political Economy* (1936), 1–30.
- Sobel, Joel, "Giving and Receiving Advice," Preprint, Department of Economics, University of California at San Diego, 2011.
- Wolitzky, Alexander, "Communication with Tokens in Repeated Games on Networks," Preprint, Stanford University and Microsoft Research, 2013.