

The Social Equilibrium of Relational Arrangements

Parikshit Ghosh

Email: Pghosh@econdse.org
Department of Economics,
Delhi School of Economics

Debraj Ray

Email: debraj.ray@nyu.edu
New York University and,
University of Warwick

[Working Paper No. 336](#)

Centre for Development Economics
Delhi School of Economics
Delhi- 110007

The Social Equilibrium of Relational Arrangements

Parikshit Ghosh (r) Debraj Ray[†]

February 2023

Abstract. The enforcement of relational contracts is especially challenging in anonymous environments when there are opportunities to start new partnerships after a transgression. Building on Ghosh and Ray (1996), we study norms within bilateral partnerships that exhibit gradually increasing cooperation, thus serving to deter deviations. However, socially beneficial gradualism may be undermined by partners renegotiating to greater cooperation from the outset. We show that incomplete information regarding partner patience ameliorates this tension even as it adds to the anonymity of the environment. Specifically, gradualism is now bilaterally desirable, and has the social by-product of maintaining individual cooperation. We also study a one-sided version of this problem in which only one of the partners exhibits moral hazard, and offer tentative thoughts on generalizing the theory to environments with richer gradations of incomplete information.

JEL Classification: C73, D85, D86.

Keywords: relational contracts, social norms, gradualism, trust-building, dynamic games.

[†]Ghosh: Delhi School of Economics, pghosh@econde.org; Ray: New York University and University of Warwick, debraj.ray@nyu.edu. Ray acknowledges funding under National Science Foundation grant SES-1851758. We thank Gerd Muehlheusser and two anonymous referees for helpful comments. “(r)” indicates author names are in random order.

1. INTRODUCTION

Many relationships — especially those with transactions separated in time or involving unverifiable actions along some dimensions — live in the shadow of ever-present opportunism. Borrowers may default on loans, workers might shirk, suppliers could cut back on quality and insurers may refuse to honor valid claims. Efficient market exchange rests on a pillar that is hidden from view in the Walrasian model — an effective institutional mechanism for contract enforcement. Consequently, cracks and weaknesses in this pillar, and possible ways to reinforce its foundations, remain beyond scrutiny in the world of classical price theory.

Anyone who has grown up in a modern, affluent society might be tempted to take a strict Hobbesian view of the world and pronounce that contracts can only be upheld by a well functioning judicial system. Indeed, a commitment to the rule of law is often taken as one of the commandments for poorer societies seeking a path to prosperity. While this dictum is undoubtedly to be taken seriously, we also realize that things are a little more nuanced when we cast our glance at the complex web of informal transactions embedded in a variety of relationships; e.g., medieval trading (Greif 1993), village moneylending (Aleem 1990), reciprocal relationships such as co-insurance (Udry 1994), immigrant networks (Munshi 2003), and so on. Commerce flourished throughout history and continues to do so in vast swathes of the developing world without an army of lawyers and judges watching over people's shoulders.

It is not a new realization among economists that contracts are often implicit and *relational* arrangements — their integrity is upheld not only by the discipline of courts but also (and sometimes exclusively) by the fact that the parties are engaged in long term business relationships which may diminish and even rupture due to a breach of promise. In this view, contracts are often *self-enforcing*, sustained by long term self interest rather than legal sanctions or ethical restraint. Indeed, even in advanced economies, deals worth billions of dollars are made as much on the basis of trust and reputation as legal safeguards (Macaulay 1963). After all, it is precisely when the carpet gets pulled out from under the investor's feet, as in the Bernie Madoff scandal, that we are reminded that it was there to begin with.

Both of us grew up in India in the sleepy days before the economic liberalization of 1991 infused new energy (or frenzy, in the eyes of some) into its corporate boards, stock markets and law offices. By that time, both of us had left India's shores for the United States, so the idea that economic transactions are highly personalized, informal and repeated, had a familiar, un-exotic ring to it, much like a comforting childhood memory. While policymakers in major international organizations agonize long and hard over critical missing markets — say, for crop insurance or institutional credit — it was not difficult to find ingenious transactions in every nook and cranny. For example, we knew (second-hand, we hasten to stress) how some enterprising “insurance agents” struck unscrupulous deals with many regular commuters who travelled on a suburban railway network without purchasing tickets. Ticketless passengers paid a monthly premium to the agents, who would reimburse them for the fines the former had to pay when the inspectors occasionally caught up with them. Such informal arrangements flourished not with the blessings of the law, but — in this instance — to arbitrage its statistical failings.

In the summer of 1993, Ghosh was a graduate student in Boston University looking for a dissertation topic, while Ray was his thesis supervisor thick in the middle of writing his textbook *Development Economics*. The fascinating quirks of factor markets in developing countries were

uppermost in our academic mindspace. In the previous two decades, contract theory customized and applied to a poorer country context had yielded rich insights into phenomena like sharecropping, informal lending and interlinkage (see, for instance, the essays in Bardhan 1989), but we wondered if decentralized arrangements in such environments could be fully understood without paying attention to the equally decentralized nature of their enforcement. Closer to home was a literature that endogenized enforcement by modeling interactions as a repeated game between a *fixed set* of players and applying it to such matters as product quality (Klein and Leffler 1981, Shapiro 1983), sovereign debt (Eaton and Gersovitz 1981), informal insurance (Coate and Ravallion 1993), ROSCAs (Besley, Coate and Loury 1994), and so on. Even more relevant was a cluster of papers (Bendor and Mookherjee 1990, Kandori 1992, Greif 1993, Greif, Milgrom and Weingast 1994, Ellison 1994, Okuno-Fujiwara and Postlewaite 1995, among many others) that studied *community enforcement* of good behavior with infractions in decentralized bilateral relationships often punished by equally decentralized third parties. These papers were freshly published at the time, or circulating as working papers. Relational contracting was very much in the air.

These issues were fascinating, to be sure, but we wished to combine them differently to address our principal concern. We were taken by the idea that as a country traverses the arc from a poor, rural society to an urbanized and industrialized modern economy, the flow of information about agents' past conduct would first worsen and then improve. To appreciate this, one need look no further than two extremes. In small village communities, disputes quickly become public information, resulting in the social ability to sustain all sorts of reciprocal or cooperative relationships. In modern market economies, there is digitized tracking of credit histories, business balance sheets and criminal records, which again leads to the ability to maintain a plethora of relationships. But in transition economies matters could well be different. Migration, mobility, the rise of new occupational sectors and the constant churning that accompanies economic development all create a cloak of anonymity well suited for fly-by-night operators. This is apparent, for example, if one contrasts Udry's (1994) study of personal loans in Nigerian villages that are enforced by community pressure, against Aleem (1990), whose investigation of credit markets in the more prosperous yet still underdeveloped Chambar district of Pakistan reveals extremely high costs of borrower background checks.

Markets in developing countries are a relational contract theorist's dream. They are often dominated by long-term relationships, with steep entry barriers into new ones (see, e.g., McMillan and Woodruff 1999). Often, there is a sharp departure from the law of one price. For instance, Aleem (1990) finds interest rates ranging from 80 to 200 percent within a single district — competitive pressures do not lead to price convergence. The duration of the relationship matters as well. Specifically, payoffs improve as the relationship progresses. One would imagine this to be arising from combination of selection and moral hazard. In the case of repeated credit, this would imply larger loans at possibly lower interest (Ghosh and Ray 2016). Several empirical studies confirm this in countries as diverse as Thailand (Siamwalla et al. 1990), Vietnam (McMillan and Woodruff 1999) and Madagascar (Fafchamps and Minten 1999). Macchiavello and Morjaria (2015) present evidence that the reputation channel is vital in such relationships. Faced with a supply shock due to political violence, rose exporters in Kenya cut back on promised deliveries, but less so the older the relationship with the client, as there was more reputational capital at stake.¹

¹Of course, as development picks up pace and legal institutions assume a more prominent role, formal and relational contracts come to coexist. In theory, their interaction could be one of complements or substitutes (Bodoh-Creed 2019) due to opposing selection and incentive effects. The presence of a sector where contracts are legally enforced raises

So much of this stems from lack of information. Clearly, the development path — from relative immobility to mature growth with all its attendant institutional underpinnings — is not a monotone one, at least as far as the flow of information is concerned. That brought us to two critical features that the literature had neglected thus far — or rather, had not considered together. This was the *voluntary continuation* of (perhaps initially random) matches and an ambient *anonymous environment*. The first feature is self-explanatory: individuals may initially be matched to their employers, creditors and vendors by chance, but they certainly *choose* whether to continue the relationship. The second feature emphasizes the idea that in anonymous environments, agents might learn about others' behavior and traits *only* through interacting with them personally. This specification entertains precisely the opposite extreme of full information flows. A new partner could be an opportunistic deviant, a harmless new entrant, or someone who has been wronged by a previous partner. This is the setting studied in Ghosh and Ray (1996), using a fairly general symmetric stage game with two-sided moral hazard.

But if deviant behavior goes undetected, what is there to prevent an anonymous environment from being overrun by unfettered opportunism, with serial defaulters, shirkers and snake oil salesmen merrily lying in wait at a revolving door of disappointed lenders, employers and clients? Surely, there must be some cost to starting new relationships. Such costs can be generated by matching frictions (e.g., the equilibrium unemployment of Shapiro and Stiglitz 1984 and Macleod and Malcolmson 1989) or perhaps costly gift exchange at the initiation of a new relationship (Carmichael and Macleod 1997). It could also arise through *gradualism*, with every new relationship expected to initialize at small levels of cooperation, generating modest payoffs, and scaling up over time conditional on past compliance. Such expectations could arise via *social norms* — behaviors most people in society are known to adhere to. Our focus was on *self-sustaining* social norms, those which give every rational agent a reason to act in accordance with the norm, as opposed to internalized values inherited from parents or society (Frank 1988, Fehr and Schurtenberger 2018). More precisely, these norms must possess a bootstrap property — deviating from it at any point implies a return to the start of the same prescribed path of actions via a new relationship, and that must be unattractive. That leads to a core idea — *the path prescribed by the norm must serve as its own punishment*. Our analysis has a general equilibrium character, where the outside option in every relationship is endogenous, unlike other papers in the literature that explored gradualism with exogenous outside options (Datta 1996, Watson 1999, 2002, Ray 2002; one exception was the contemporaneous work of Kranton 1996).

There are many self-sustaining norms, including the trivial one where every agent always plays the one-shot best response and there is no reason for long term partnerships to form. What is of interest is an *efficient* self-sustaining norm, one that creates the highest payoff in the class of all self-sustaining norms. In Proposition 1, we identify the “payoff destruction” that must occur under an efficient norm whenever a relationship has to be restarted. Proposition 2 shows there is a gradualist path that can achieve this — no extraneous device like gifts or frictions are necessary. Neither the individual nor society will want to deviate from this path under the circumstances.

This could be a tidy end to the story until one realizes that deviation from a norm could occur at three levels — the individual, the community and the relationship. Having some cost attached

outside options of agents engaged in informal business relations, but could also draw away untrustworthy agents from their midst, improving the pool of partners. Johnson, McMillan and Woodruff (2002) find evidence of complementarity in post-Soviet economies — as courts became more reliable, trust in business partners went up. The interaction between the two kinds of institutions merits further theoretical and empirical study.

to the beginning of any relationship may serve a useful social role but any specific relationship would want to do away with it and jump to full-blown cooperation right away as long as the front-loaded cost in *other* potential relationships serves to reduce the value of their outside option. This, however, undermines and eventually destroys the gradualist norm itself.

To repair this cleavage, and taking a cue from the reputation literature (Kreps et al. 1982), we introduced a second dimension of private information — one that has to do with agents' inherent traits — by assuming that a fraction of the population are habitual cheats who always play the short term best response (they can be thought of as behavioral types or short-lived or myopic agents). This provides a rationale for gradualism even at the level of a single relationship. Agents who are interested in building long term relationships may want to start one by proposing a small, experimental level of cooperation that does not put too much at stake but nevertheless separates the shortsighted cheats from farsighted agents. For example, lenders could give a “testing loan” (Aleem 1990) to a new client or buyers could place a small order from a new vendor to verify trustworthiness. In Ghosh and Ray (1996), we demonstrated the existence of such trust-building equilibria that are renegotiation-proof at the level of the bilateral relationship — a condition we called *bilateral rationality*.

Thus a problem created by one kind of informational friction (the absence of information about past actions) can be at least partly ameliorated by *additional* information frictions (the absence of information about agent types). Proposition 3 in this paper goes beyond the Ghosh-Ray analysis to ask a deeper question about this serendipitous interplay. Notice that the relationship-specific conservatism at the initial stages occurs only to avoid large losses from the behavior of a potentially opportunistic partner. How well does this conservatism work to shore up the social environment? Surprisingly, as long as there are enough myopic types in the environment to permit the existence of a social equilibrium, no conflict arises between relational and social goals — the bilateral rationality constraint does well in generating an efficient self-sustaining norm. Of course, the lack of formal enforcement and a public repository of information about individual histories are both costly. Nevertheless, in societies lacking such state capacity, the best constrained form of economic activity can arise in a *decentralized* way, as agents try to forge the best possible economic relationships they can build given their social environment. This lends an “invisible hand” flavor to our results, but in a third-best world with at least two kinds of informational frictions.

This paper is written as a primer for the particular sub-area of “the general equilibrium of relational contracting.” We summarize some of the contributions in this area, and use our own paper (Ghosh and Ray 1996) to organize the material. Our goal is to provide an expository account of existing findings, but going beyond that, to engage in a discussion of some of the conceptual issues that are involved in extending and generalizing already-available results, for example to a continuum of types or one-sided moral hazard problems with transferable utility.

2. MATCHING WITH PERSISTENT RELATIONSHIPS

2.1. Baseline Model. Pairs of players are drawn from a continuum population, and matched to play a stage game. While an initial match is randomly drawn, it can be continued into the indefinite future if the matched partners so wish. Specifically, at any date, after the stage game is played, each player in any pair has the option to continue the previous match or to terminate it. If both players agree to continue, they play the stage game between themselves once more. If at least

one player decides to terminate, then their relationship is broken, and both players return to the unmatched pool. As in Jackson and Wolinsky’s (1996) solution concept for link-formation in networks, we can avoid trivial coordination failures in these continuation decisions by assuming that an existing match is renewed if both players see a mutually advantageous continuation. At the next date, the same story is repeated. Meanwhile, the pool is exogenously rejuvenated by a steady inflow of new individuals and/or by exogenous player separation, the probability of which will be folded into their discount factor.

A stage game has the standard format, but for the purpose of this paper we express it in starkly symmetric form. A level of cooperation x is proposed. Implementing x will generally require a pair of actions to be undertaken, one by each player, but we black-box that entire process here for expositional ease.

Let $v(x)$ denote the common stage payoff when both players play the action x . Let $g(x)$ be the best possible deviation gain to a player when her partner sticks to her end of the cooperative bargain proxied by x . (So the corresponding deviation payoff is $v(x) + g(x)$.) We assume that $v(x)$ is single peaked and strictly increasing from 0 till $x^1 = \max v(x)$. Without loss of generality, in what follows, we restrict attention to the range $[0, x^1]$, where cooperation is valuable. Over this range, assume $g(x)$ is continuous and increasing, with $g(x) > 0$ for all $x > 0$, and with $g(0) = v(0) = 0$. When a player deviates to her best gain, the non-deviating player earns a “loser payoff” $l(x) < v(0) = 0$, which we take to be continuous and decreasing in x . It should be clear that the unique equilibrium in a one-shot version of this game is one in which both individuals engage in zero cooperation.

Though we discuss matching frictions below, we emphasize the idea that a new matching can be easily found. For instance, we can presume that there is always an exogenous rate of breakup of existing matches, even when deviations have not taken place. Or perhaps there is always a constant flow of new entrants into the pool(s). Then a deviator, who is just an atomless point in a continuum, can always re-match with ease. In short, we have removed the possibility of non-matching by assumption. In such situations, it becomes obvious that all deterrence to deviations will have to emerge from the path of the relationship itself.

2.2. Some Remarks on the Setting. Our model captures certain salient features of interest. First, it is rarely the case that repeated interaction occurs across a fixed set of partners. Nor is it necessarily the case that partners are always re-matched after every interaction. Most ongoing economic interactions arise as voluntary partnerships, as in a marriage or an employer-employee relationship, where the initial meeting is indeed often random, or is driven to some degree by stochastic elements. But after that initial meeting, the relationship can be continued, typically with the consent of both parties, and sometimes broken. That is the general background we seek to describe.

Second, we are interested in the maintenance of cooperative “norms” in the complete absence of public information flow about histories. When a matched player deviates and re-enters the pool of unmatched agents, she is indistinguishable from an agent newly-arrived to that pool, or from an agent who has been separated for exogenous reasons not involving a deviation, or indeed from an agent who has been separated from her partner owing to the partner’s deviation. This is not to say that such public records cannot be modeled in this setting. As in Rosenthal and Landau (1979), Bhaskar and Thomas (2019) and Clark, Fudenberg and Wolitsky (2021), we could presume that

each player in a newly matched pair arrives with a record that depends on their past history of play. But our focus is emphatically on the case in which no such past history exists.

Third, while a deviation can potentially be punished *within* a fixed match as in a repeated game, there is little to be gained in doing so. No punishment payoff can be pushed below the continuation value of a new relationship, because the deviating player can unilaterally exit the relationship. Therefore, at least in the realm of subgame perfect equilibria without further restrictions, the repercussions of a deviation can be viewed as tantamount to a subsequent severance of the relationship.

Fourth, we presume that agents are matched in a symmetric partnership, with moral hazard equally present on both sides of the relationship. This fully symmetric “two-sided model” is explored by Ghosh and Ray (1996), Kranton (1996), Watson (1999), Lindsey, Polak and Zeckhauser (2000), Eeckhout (2006) and Fieler (2007), among others. The leading example here would be an equal partnership, either of the social or the economic variety. But one can also conceive of a “one-sided” variant, in which players are matched from two pools (say “principal” and “agent” pools). The subsequent relationship could permit the principal to offer a contract with commitment within the period, while the agent honors or breaks the terms of that contract. For instance, Datta (1996) and Wei (2018) work in a credit-market setting, in which the lender advances a loan and the borrower can deviate by defaulting on that loan. Or one can accommodate employer-employee applications as in Eswaran and Kotwal (1985) and Shapiro and Stiglitz (1984). The variable x would then have to stand in for an index that varies both the size of the loan and the interest rate, or hours worked and the wage rate, in a way that benefits both parties. It is also possible to directly treat x as a multidimensional variable, perhaps incorporating transfers with opposing implications for payoffs to each party; see, for instance, Ghosh and Ray (2016). We return to this extension in Section 6.2.

2.3. A Remark on Cooperation Via Contagion. An alternative approach is that behavior in any one relationship affects the entire social norm in a way that is internalized by every individual. This is a face-value rendition of what our parents told us: “behave well, for your bad deeds will come back to haunt you.” Cooperation can then be achieved if individual players are patient enough. In large finite populations, the notion of contagion implicit in that parental admonition finds a channel through the game-theoretic lines of the folk theorem. Kandori (1992) and Ellison (1994) show that possibilities for cooperative play exist for finite populations if the discount factor is sufficiently close to 1. For concreteness, think of the standard two-action prisoner’s dilemma in a random-matching environment without endogenous continuation, as in Rosenthal (1979). Suppose that all players are initially matched in pairs, and begin by playing the “cooperate” action. They continue to do so until they meet a player who plays “defect,” and themselves switch to “defect” thereafter. Bad behavior, once initiated, will then spread in the population following the dynamics of an infectious disease. Well-known incentive constraints need to be satisfied for such contagion strategies to constitute an equilibrium.² Then, given any large (but finite) population, there is some threshold discount factor above which a potential deviator will pause — faced with the possibility of sufficient punishment engendered by the contagious nature of his actions. As Ellison (1994) shows that contagion can be remarkably speeded up, depending on the locally interactive structure of the matching process. Deb (2020) and Deb, Sugaya and Wolitsky (2020) extend these ideas in innovative and powerful ways to generate folk theorems for anonymous random matching games.

²An individual must find it in her interest to not start a contagion in the first place, but should also not want to slow down a contagion (by continuing to cooperate) after detecting a deviation.

Exciting though these contributions might be from a game-theoretic perspective, it is not the line adopted here. With full respect for the “contagion-based” literature, we believe that the prospects along these lines are dim for agents attempting to forge efficient transactions in large anonymous populations. (For smaller populations or with interactions that are largely local, the contagion arguments may well make good sense.) This is why we presume that the population is infinite — effectively, a continuum. The researcher must decide which model to adopt in particular situations.

In summary, the key to forming long-term cooperative relationships in our setting is either to have gradualism in those relationships, or some form of “entry barrier” into new relationships. These are the two avenues of punishment when public histories are missing. In contrast to the contagion literature, we start from the premise that the mere threat of termination of the ongoing partnership is not enough if the player who cheated can easily start another cooperative partnership.

3. SUSTAINABLE NORMS

3.1. Baseline Concepts. A *norm* is a path of prescribed actions or cooperation levels $\mathbf{x} = \{x_t\}_{t=0}^{\infty}$ taken over the full course of a matched relationship, presuming that it lasts. Let V_t be the associated sequence of continuation values generated by the norm; that is,

$$(1) \quad V_t = (1 - \delta) \sum_{s=t}^{\infty} \delta^{s-t} v(x_s)$$

where we follow the standard practice of normalizing by $1 - \delta$ so that these lifetime payoffs are “in the same units” as their one-shot counterparts. Given our working hypothesis that the forming of a new relationship occurs with no difficulty, we are led naturally to the following definition: a norm is *sustainable* if for every date t ,

$$(2) \quad \delta[V_{t+1} - V_t] \geq (1 - \delta)g(x_t).$$

That is, an individual i can deviate on an ongoing \mathbf{x} at any date t . She does so by enjoying her deviation gain $g(x_t)$. Thereafter she is back in the common pool and gets the payoff V_0 which accrues at the beginning of a new relationship. Condition (2) states that the one-shot payoff gain from such a deviation should not exceed the future benefits of continuing on the path \mathbf{x} instead of returning to its initial point. The bootstrapping implicit in (2) is a stark consequence of costless rematching, a feature we have chosen to highlight.

We might additionally ask for a sustainable norm to be efficient. That is, \mathbf{x} maximizes

$$(3) \quad V_0 = (1 - \delta) \sum_{t=0}^{\infty} \delta^t v(x_t)$$

in the class of all sustainable norms.

3.2. Remarks on Sustainability. Some points of interest are to be noted. First, the continuation values $\{V_t\}$ are defined on the presumption that the norm is sustainable to begin with. Is there any contradiction or circularity here? The answer is no, as is well-known to anyone familiar with the one-shot deviation principle in discounted dynamic programming.

Second, we have presumed that the norm \mathbf{x} lasts into the indefinite future. What if it is destined to end at some finite date — that is, the norm includes the inevitability of separation after, say,

T periods of interaction? It is then easy to see from a backwards induction argument, and our assumption that the stage game has a unique Nash equilibrium, that such a norm must involve the play of that static equilibrium at every date.

Third, we have presumed that there is a single norm of play. What if there are several — call them x^1, \dots, x^m ? Then for each potential deviator in some current relationship x^k at date t , there is some probability of accessing one of the different going norms following a deviation, possibly mediated by observable characteristics of the individual: gender, race, and so on. The resulting incentive constraints can then be written as a straightforward extension of (2).

Fourth, it is easy to accommodate a cost of rematching; simply subtract that additional cost from the left-hand side of (2). However, those costs could be endogenous to the norm itself, as in Ghosh and Ray (2016) and then will require more care to incorporate; we discuss this briefly below.

Finally, notice how the definition of an efficient sustainable norm highlights the heart of the no-information problem. Such a norm seeks to maximize V_0 , but V_0 is also the punishment value that holds that very norm together. To create stronger incentives for cooperation, it is helpful to reduce the value of a player's outside option, which is the lifetime payoff at the beginning of a new partnership. However, maximizing the value of partnerships is the assumed teleology of the norm. How best to handle this tension is the heart of the matter.

Datta (1996), Ghosh and Ray (1996, 2016), Kranton (1996), Watson (1999, 2002), Wei (2018) and Watson and Hua (2022), among others, explore the idea of *gradual* trust-building and its effect on disciplining relationships. While these papers focus on different aspects of gradualism, there is a common core to the arguments that we extract and highlight in this section. But more importantly, our focus is on the social or economy-wide repercussions of such interactions.

3.3. A Simple Norm When Payoff Destruction is Possible. To gain some insight into the nature of the problem, consider a closely related exercise. Suppose there is some instrument that allows any part of the expected payoff to be destroyed at the beginning of a partnership (some examples will follow soon). Let $d \geq 0$ be the magnitude of this initial payoff destruction, expressed in utils. Let $V_t(x)$ be the lifetime continuation payoff generated by a path x at date t , as given by (1). After applying payoff destruction at the beginning of a partnership, we have $V_0 = V_0(x) - d$ and $V_t = V_t(x)$ for all $t > 0$.

Now consider maximizing V_0 by choosing a pair (x, d) so that the associated values as just defined satisfy the no-deviation constraint (2) at all dates. This problem has a simple solution. Define

$$(4) \quad V^* = \max_x \left[v(x) - \frac{1 - \delta}{\delta} g(x) \right],$$

and let X^* be the associated set of maximizers. We assume throughout that $V^* > 0$.

PROPOSITION 1. *Under an efficient sustainable norm, $V_0 = V^*$. For any $x^* \in X^*$, one such norm is given by $x = (x^*, x^*, x^*, \dots)$ and $d = v(x^*) - V^*$.*

With unlimited capacity for payoff destruction, x^* is the highest level of cooperation reached by an efficient sustainable norm. However, without such destruction, the maintenance of x^* (or any fixed positive level of cooperation for that matter) from the beginning of every partnership is problematic, since partnerships can be broken and restarted costlessly. That leads to the general idea

that there must be “enough payoff destruction” whenever a new partnership is started. What are plausible ways in which such payoff destruction may come about? The first two instances that now follow are not our central focus, but they serve to place our model in context. The third is central.

3.4. Money Burning and Gift Exchange. A straightforward path to payoff destruction is the incorporation of some obligation to burn money at the beginning of a new relationship. The optimal level of money burning which maximizes the value of a partnership at its starting point is given by its utility equivalent $d = v(x^*) - V^*$, as in Proposition 1. Carmichael and MacLeod (1997) analyze the widespread practice of gift exchange at the beginning of new relationships — expensive engagement rings, lavish business dinners, etc. Such gifts must be intrinsically inefficient — they should have lower use value to the recipient in comparison to what they cost to the giver. Inefficient gifts make starting new relationships costly and serve the strategic purpose of money burning.

3.5. Matching Frictions. Starting a new relationship could be time-consuming or expensive. Such matching frictions may be introduced exogenously into the model (as in Greif 1993) or they could arise endogenously; e.g., via equilibrium unemployment in labor markets (Shapiro and Stiglitz 1984, Eswaran and Kotwal 1985). Suppose that an unmatched agent meets a new partner with probability α in any period. Given a norm x , let \tilde{V}_0 be an individual’s lifetime expected payoff at the start of a relationship *once a partner has been found*, and V_0 her lifetime payoff at the start of a search. Then $V_0 = \alpha\tilde{V}_0 + (1 - \alpha)\delta V_0$, which implies

$$(5) \quad V_0 = \left[\frac{\alpha}{1 - \delta(1 - \alpha)} \right] \tilde{V}_0.$$

If $x = (x, x, \dots)$ is stationary, then $\tilde{V}_0 = v(x)$ and $V_t = v(x)$ for all $t \geq 1$. Then for any $\alpha \in [0, 1]$, the maximal level of cooperation $x(\alpha)$ is obtained by combining (5) and the binding incentive constraint (2), and is given by the largest solution (in x) to

$$(6) \quad \delta(1 - \alpha)v(x) = [1 - \delta(1 - \alpha)]g(x).$$

If α is endogenous, as in the equilibrium unemployment model, a separate condition — possibly involving x — will serve to determine it. That condition, along with (6), will then pin down both x and α . In general, high values of α lead to lower cooperation: for instance, as $\alpha \rightarrow 1$, $x \rightarrow 0$.

Matching frictions might arise out of impediments in the search process, as already mentioned, but could also have sociological roots. For instance, cooperative norms may be identity-based, with long term cooperative relationships only with in-group members and eschewal of such attempts with outsiders (Ruffle and Sosis 2006). Endogamous economic relationships (e.g., ethnic networks in job seeking as in Munshi 2003), or mutual insurance in caste networks (Munshi and Rosenzweig 2009) are widely observed. Such identity based segregation creates frictions in the path of forming cooperative relationships, thus — paradoxically — helping their ultimate formation.

3.6. Gradualism. Could a norm be self-sustaining without auxiliary devices such as gift exchange or matching frictions? The answer is in the affirmative, but such a norm must be carefully constructed. (For instance, a constant level of cooperation is clearly unsustainable.) The following observation shows that there exist sustainable norms that attain the theoretical upper bound on starting payoffs made possible by payoff destruction or matching frictions. The path x itself can serve as a substitute for such devices.

PROPOSITION 2. *Under an efficient sustainable norm \mathbf{x} , $V_0 = V^*$, and $\lim_t x_t = x^*$ for some $x^* \in X^*$. For any $x^* \in X^*$, one efficient sustainable norm which transitions to x^* in the shortest time takes the following form:*

$$(7) \quad \mathbf{x} = (\underbrace{0, \dots, 0}_T, \hat{x}, \underbrace{x^*, x^*, \dots})$$

where T is the smallest integer such that $\delta^T v(x^*) < V^*$ and \hat{x} satisfies $\delta^{T-1}(1-\delta)v(\hat{x}) + \delta^T v(x^*) = V^*$

Proposition 2 states that every efficient sustainable norm must converge to some maximal cooperation level x^* in finite time. That level is preceded by an initial trust-building phase where cooperation levels are strictly lower. Such backloading of cooperation and high payoffs, or *gradualism*, is an intrinsic feature of efficient norms in this environment. Without it, there would be no cost to the termination of a partnership and consequently no incentive to refrain from cheating.

Gradualism can also arise in repeated games between a fixed set of partners both on-path (Ray 2002) as well as off-path (Abreu 1988). However, it only arises for a certain class of stage games; for others, optimal play as well as optimal punishments are stationary.³ The gradualism that arises in this environment is *inevitable* if non-trivial levels of cooperation are to be achieved. Repeated games with imperfect monitoring (Abreu, Pearce and Stacchetti 1990) do share the feature that cooperative and punishment paths may coincide on their tails.

The quickest path to cooperation involves an initial waiting period when no cooperation is attempted (see also Kranton 1996 and Fujiwara-Greve and Okuno-Fujiwara 2009). But the initial trust-building phase is not uniquely pinned down in general. Fielor (2007) emphasizes the three-phase structure in (7), though there may be other efficient norms that do not display this property. But all efficient norms must eventually reach the last steady-state phase of maximal cooperation. Under additional assumptions, Datta (1996) establishes continuation payoffs must be non-decreasing along the path of play, and Wei (2018) shows that the same is true of cooperation levels, under any efficient norm.

It is instructive to compare x^* against some benchmarks. The *first-best* cooperation level x^1 is the value of x which maximizes $v(x)$. In a world where contracts are enforced by a legal authority, this is the cooperation level that will be agreed upon from the outset. Next, define the *second-best* cooperation level x^2 as the highest stationary value of x that maximizes $v(x)$ subject to the no-deviation constraint when $V_0 = 0$, i.e., under the constraint $\delta v(x) \geq (1-\delta)g(x)$. It is the cooperation level that would obtain in a world of relational self-enforcing contracts, but one which allows for maximal punishment. This can arise, for example, if partner switching opportunities are altogether absent, or if player histories are publicly observed and there is adequate scope for community enforcement (Kandori 1992). Since x^1 is the unconstrained maximand, we must have $v(x^2) \leq v(x^1)$. Finally, note that x^* satisfies the constraint of the maximization problem that yields x^2 as its maximand, since from (4), $V^* = v(x^*) - \frac{(1-\delta)}{\delta}g(x^*) \geq 0$. This implies $V^* = v(x^*) \leq v(x^2) \leq v(x^1)$, with the last inequality strict under suitable curvature assumptions.⁴

³For the stage game we consider here, a player's minmax payoff coincides with her one-shot Nash equilibrium payoff. Consequently, in the repeated game between two fixed players, the optimal punishment path is indeed stationary — it involves perpetual repetition of the one-shot Nash equilibrium play.

⁴For example, assume $v(x)$ is strictly concave and $g(x)$ is strictly convex. Then, $v(x) - \frac{(1-\delta)}{\delta}g(x)$ is strictly concave and attains its maximum value (at x^*) strictly before it hits 0 (at x^2).

Note that cooperation levels that exceed x^* and up to the second best x^2 can be *eventually* reached along *some* self-sustaining path. But these higher levels of cooperation will require a longer waiting time T for which cooperation remains suspended at the start of a relationship. Such paths are sustainable, but not efficient. Greater cooperation must be traded off against the longer time it takes to get there — a trade-off which is typically resolved at some interior point. The important takeaway is this: the requirement that contracts be self-enforcing carries a welfare cost, and the need to make them work in *anonymous* environments adds to that cost.

4. CAN EFFICIENT NORMS BE DECENTRALIZED?

Imagine that an efficient norm is in place, and then matched pairs “best-respond” to find the best path that they could then support. Is the norm robust to such best responses, or is it likely to be undermined by partnerships that find it profitable to move away from the prescriptions of the norm?

4.1. The Shadow of Renegotiation. Sustainable norms are certainly self-enforcing in the sense that no *single* agent can profit from a unilateral deviation, as long as everyone else in society adheres to the norm. However, matched pairs of agents might have access to alternative arrangements that are in the interest of both parties (and which are also immune to unilateral deviations), *taking as given* the custom followed in the rest of society. After all, while norms might have value in reducing the scope for deviation at the social level, that generally comes with reduced efficiency in individual partnerships. There is no reason why such partnerships would bear that inefficiency for social purposes.

For instance, money burning may have excellent social value, but as long as it is practiced elsewhere, two players starting a new relationship have good reason to drop this costly ritual for themselves. Similarly, it makes sense for any given pair to cut the initial waiting period (as described in Proposition 2) and jump to maximal cooperation right away. In other words, since efficient social norms involve payoff destruction in some form or the other (Proposition 1), it is vulnerable to pairwise deviations even if it is protected against individual opportunism. It would appear that cooperation in anonymous environments will collapse under the weight of this contradiction.

4.2. Incentivized Caution Via Incomplete Information. Putting this last observation another way, *caution* performs the social function of shoring up relationships, because a price has then to be paid for entering a new match. But that caution needs to be internalized in individual interactions. In Ghosh and Ray (1996), we made the argument that the existence of incomplete information across matched players creates grounds for caution even after bilateral rationality is imposed. That caution — or deliberate gradualism — will then seep into the social norm and could restore its sustainability. In this way, missing information about player *types* can help undo the damage caused by missing information about player *histories*, which is another one of those paradoxical but presumably useful messages from the land of the second-best.

We formulate this argument by introducing heterogeneity in individual farsightedness. Assume, then, that in the unmatched pool, there is a probability measure π_0 over discount factors.⁵ Each partner in a match knows her own type but has initial belief π_0 over the type of her partner. These

⁵This measure may itself be endogenous to social behavior (see Ghosh and Ray 1996 for a discussion) but in the interest of expositional clarity we presume that π is given.

beliefs will evolve over the course of the relationship, but given the simplicity of the relationship (abide, defect), any agent’s belief π about the partner at any stage will be the same as the partner’s reciprocal belief about *her*. A social norm \mathbf{x} is now more than just a path: it prescribes a level of cooperation $x_t(\pi)$ for each belief π and stage t , and it specifies exit decisions $d_t(\pi, \delta) \in \{0, 1\}$ for each stage, belief and each δ -type, where $d = 1$ means that the individual takes her stage-dominant action and restarts a relationship with a newly matched partner.⁶ The collection of bilateral and individual plans prescribed by the norm will — if universally adhered to — “unfold” into an outcome path as π evolves in accordance with Bayes’ Rule.

This definition of a norm handles all conceivable beliefs at all conceivable stages. However, with only individual and bilateral deviations to be considered, and a continuum of individuals adhering to the norm, all that matters are *on-path beliefs* — that special sequence π_t generated by adherence to the exit recommendations of the norm, starting from π_0 .⁷

We are interested in efficient, sustainable norms that meet both individual and bilateral requirements of incentive-compatibility. Such norms will be able to select outcomes over and above those generated by individual best responses alone. For instance, efficient sustainable norms might avoid trivial coordination failures where it is believed in a self-sustaining way that every type will defect. In part, norms attempt to avoid these failures, but of course they do much more, including the testing and weeding-out of relatively impatient types as the relationship proceeds.

4.3. A Simple Two-Type Setting. Suppose that in the unmatched pool, a fraction π of agents are forward-looking or patient players, equipped with some discount factor $\delta \in (0, 1)$. The remaining fraction, $1 - \pi$, consists entirely of myopic or impatient players who maximize their current payoff in the stage game. Given our assumptions, this means they always choose the dominant action 0. In this simple setting, a sufficient test to weed out an impatient type is the offering of *any* strictly positive level of cooperation. If an individual passes that test, she is deemed to be patient.

An efficient norm will maximize expected payoffs subject to no-deviation constraints; the latter are only relevant to patient types as just noted. Note that “maximum expected payoffs” is an ambiguous notion. It depends on the “welfare weights” assigned to patient types. Temporarily sidestepping this issue, we presume that society — and every matched pair — only values the payoff to the patient type, but see Section 5 below, and especially Sections 5.2 and 5.3.

As for the structure of the norm, there are only two on-path beliefs that matter: $\pi_0 = (\pi, 1 - \pi)$ and $\pi_1 = (1, 0)$, the former at the initial stage and the latter at all later stages (assuming that the norm prescribes nonzero cooperation). A norm \mathbf{x} prescribes an initial cooperation level x_0 at π_0 , and a sequence $\{x_t(\pi_1)\}$ of cooperation levels at every stage t thereafter. We suppress the obvious exit recommendations for expositional ease, but we know what they are: the patient type is asked to always stay, and the impatient type is asked to leave at stage 0. So the lifetime value for patient types at the start of the relationship is the solution in $V_0^{\mathbf{x}}(\delta)$ to:

$$(8) \quad V_0^{\mathbf{x}}(\delta) = \pi_t[(1 - \delta)v(x_0) + \delta V_1^{\mathbf{x}}(\delta)] + (1 - \pi_t)[(1 - \delta)l(x_0) + \delta V_0^{\mathbf{x}}(\delta)]$$

⁶This presumption raises other questions about bilateral rationality, not when things are going well, but when they are going badly. Suppose that A has deviated in her arrangement with B , but A has still proved herself to have a high discount factor. Might that not serve as grounds for a perhaps reluctant reconciliation, given that A has a discount factor that is demonstrably higher than that of the average pool? We avoid this question here.

⁷However, exit recommendations do need to be made for all δ , whether or not they are in the support of π . Individual types can entertain counterfactual scenarios in which she was “asked” to exit earlier, but did not.

where the continuation value $V_1^{\mathbf{x}}(\delta)$ is uniquely pinned down by $\{x_t(\boldsymbol{\pi}_1)\}$.

Now we turn to the bilateral rationality of the norm. At any stage $t \geq 1$, there is just the patient type left, and we want there to be no group-level improvement on the norm by some alternative proposal that respects the social norm as an outside option. So bilateral rationality demands that cooperation should be at the *highest level* y satisfying

$$(9) \quad (1 - \delta)g(y) \leq \delta [v(y) - V_0^{\mathbf{x}}(\delta)].$$

at all stages after 0. As for stage 0, the matched pair again takes the norm \mathbf{x} as given. The pair knows it is in its own interest to run a test. It will seek the maximum payoff from doing so subject to no-deviation conditions for the patient types.⁸ In short, the pair will choose an initial test level $y_0 > 0$ to maximize $\pi v(y_0) + (1 - \pi)l(y_0)$, subject to the sustainability restriction:

$$(10) \quad (1 - \delta)[\pi g(y_0) - (1 - \pi)l(y_0)] \leq \delta \pi [V_1^{\mathbf{x}}(\delta) - V_0^{\mathbf{x}}(\delta)].$$

The objective function is expected payoff from cooperation y_0 , given the current uncertainty about a new partner's type. The left-hand-side of the constraint is the one-shot payoff to a patient player who deviates from y_0 by playing the dominant action. It is a weighted average of the gain $g(y_0)$ (in case her partner is the patient type) and the avoided loss $l(y_0)$ (in case her partner is the impatient type). The expression on the right-hand-side captures the expected future loss from such a deviation, as the agent foregoes the opportunity to move on to the next stage of the path instead of restarting at stage 0. Of course, that next stage is only relevant if the current partner is the patient type, which explains the probability π on the right hand side of (10).

We pause to take stock of and eliminate a minor technical issue. The above expected payoff to (the patient types in) the pair can generally be taken to be positive by suitable choice of x , provided that mild restrictions apply — such as end-point Inada-like conditions on the functions v or l , or a suitable lower bound on π . But if anyway the payoff is unavoidably negative for all x , we adopt the convention that the test consists of extending an infinitesimal degree of cooperation (x is “tiny”) at roughly zero cost, and this will suffice to weed out the impatient type. With this convention in hand, the Stage-0 maximization problem above always has a well-defined solution.

Say that a social norm \mathbf{x} is bilaterally rational if every matched pair willingly adopts — in the sense just described — the actions prescribed by the norm at every stage. Our discussion above makes it clear that with just two types, any bilaterally rational norm should be time-stationary from stage 1 onward, at the highest level x satisfying (9). At Stage 0, a more subtle departure from social efficiency is potentially possible. The *social* construction of caution relies on the twin objective of higher welfare for a matched pair, while at the same time preserving enough ammunition for punishment in the event of a deviation. The *private* construction of caution places value only on the first of these two objectives, because it takes the ambient norm as given — a matched pair is too small to influence the entire norm. And yet, we have:

PROPOSITION 3. *Consider the two-type setting with a proportion $\pi \in (0, 1)$ of types with a common discount factor $\delta \in (0, 1)$, while the remainder have a discount factor of zero. Then:*

(i) *Presuming it exists, a bilaterally rational norm is always socially efficient.*

⁸These are delicate matters: an impatient type lurking within the pair might want to run a very different “test” that it can suitably exploit by a deviation. But such a suggestion will reveal the player to be impatient. We therefore take it that — patient or not — the pair can only run a test to maximize the payoff to the patient types. We return to this point in Section 5.

- (ii) *There is a threshold $\hat{\pi} \in (0, 1)$ such that a bilaterally rational norm exists if and only if $\pi < \hat{\pi}$.*
- (iii) *Under any bilaterally rational norm \mathbf{x} , $V_0^{\mathbf{x}}(\delta) \leq V^*$, where V^* is defined in (4).*

To see part (i), suppose on the contrary that \mathbf{x} is not socially efficient. Then there is sustainable \mathbf{z} with $V_0^{\mathbf{z}}(\delta) > V_0^{\mathbf{x}}(\delta)$. Because punishments are weaker under the \mathbf{z} -norm, it is easy to see that $V_1^{\mathbf{z}}(\delta) \leq V_1^{\mathbf{x}}(\delta)$. But then a matched pair under the ambient norm \mathbf{x} can do as well or even better than the actions along \mathbf{z} . After all, (9) is a weaker constraint under \mathbf{x} than under \mathbf{z} , because $V_0^{\mathbf{z}}(\delta) > V_0^{\mathbf{x}}(\delta)$, and so is (10), because $V_0^{\mathbf{z}}(\delta) > V_0^{\mathbf{x}}(\delta)$ and $V_1^{\mathbf{z}}(\delta) \leq V_1^{\mathbf{x}}(\delta)$, as we've just argued. That means (y_0, y) does at least as well as $V_0^{\mathbf{z}}(\delta)$ for the matched pair. Given that $V_0^{\mathbf{z}}(\delta) > V_0^{\mathbf{x}}(\delta)$, \mathbf{x} cannot be bilaterally rational, a contradiction.

This argument shows that despite the externality alluded to just before the statement of Proposition 3, a bilaterally rational norm is always socially optimal, *provided it exists*. That is, if a social planner can improve on a norm, so can a matched pair, even if it does not internalize the social externality. However, the problem is that a bilaterally rational norm may not exist, Part (ii) of the Proposition states that this happens *only* when the proportion of patient types is “too large,” and in the process it also nests our argument for the complete-information model; see Section 4.1.

So the proposition highlights the critical importance of a sufficient number of impatient agents inhabiting the unmatched pool. Given the presence of such individuals, their more patient counterparts will face (stochastic) delay in finding a partner who is willing to cooperate. This delay acts like a tax on the payoff that a voluntary long term relationship generates. Consequently, it serves the purpose of payoff destruction to some extent — but to what extent? Patient agents will test the waters at the start of a relationship by proposing an experimental level of cooperation. However, if most agents in that pool are patient, they will be emboldened to experiment with a high level of cooperation. That destroys the punishment on which an incentive-compatible norm must rest. The delay caused by accidental initial encounters with opportunists goes some way to resolve this problem, but if the proportion of myopic types is small, it cannot be enough. The efficient norm will somehow have to build in an additional waiting phase, but that wait is not bilaterally rational.

As the proportion of impatient agents grows, two things happen. First, owing to the increased risk of being cheated, patient agents willingly lower the level of desired cooperation at the beginning of a relationship. Second, the need to build payoff destruction into the path \mathbf{x} itself is lower, because the increased risk acts as a natural tax on payoffs. The wedge between social priorities and the bilaterally selfish interests of any given partnership is therefore attenuated, progressively so as the proportion of impatient agents grows larger. Indeed, there is a threshold fraction of impatient types for which the wedge disappears completely.

Additionally, once the fraction of impatient agents in the unmatched pool crosses this just-described threshold, the expected payoff to any patient type from an efficient norm must fall below V^* . In this situation, constrained efficiency dictates (incentive compatible) cooperation possibilities at every stage be exploited to the hilt, i.e., there must be no slack which individual pairs could subsequently exploit through renegotiation. This is part (iii) of the Proposition.

Ghosh and Ray (1996) characterized sustainable norms satisfying the bilateral rationality condition. The observation of interest in Proposition 3 is that there are many situations where renegotiation does *not* erode norms. No additional social cost need be imposed on deviators beyond what

arises from bilaterally rational calculation. In other words, the best arrangement for society can sometimes emerge from individuals trying to forge the best possible relationships.

4.4. Scarcity and Gradualism: A Remark. There are two distinct aspects that work in tandem to discipline opportunism — the *scarcity* of patient types and the *gradualism* built into the norm. Scarcity arises because not all agents are patient; gradualism is a product of the uncertainty that this creates. In our setting, scarcity is exogenously assumed in the form of π . Is gradualism additionally needed under an efficient norm, in the sense that initial cooperation must be strictly lower?

Suppose that $v(x) = x$, $g(x) = \beta x$ and $l(x) = -\gamma x$, where $\beta > 0$ and $\gamma > 0$. First, assume that $x \in \{0, 1\}$, i.e., there is only one level of positive cooperation. Some bounds must be placed on δ and π to get any cooperation from patient agents. The condition $V^* > 0$ means that $\delta > \beta/(1 + \beta)$.⁹ Second, π must not be so high that a patient individual would prefer to perpetually cheat strangers rather than enter into a cooperative relationship with the patient types among them. That is, $\pi(1 + \beta) < 1$, an upper threshold that plays the role of $\hat{\pi}$ in the previous section.

To separate patient agents from impatient ones, and given just two actions, it is necessary (though potentially hazardous) to choose $x_0 = 1$ in the initial phase. This is followed by full-blown cooperation thereafter if the partner is also cooperative: $x_1 = x_2 = \dots = 1$. Given $\beta < \gamma$, patient players will be most tempted to defect in the initial phase. The no-deviation constraint is satisfied at all stages if and only if:

$$(11) \quad \delta \geq \frac{\pi\beta + (1 - \pi)\gamma}{(1 - \pi)[1 + (2 - \pi)\gamma + \pi\beta]}$$

Return now to the case of continuous cooperation levels with $x \in [0, 1]$. What parametric restrictions are needed to sustain a path $x = (x, 1, 1, 1, \dots)$ for some x ?

The answer is that *no additional conditions are needed*; that is, (11) becomes redundant. If (11) is satisfied, well and good, but if not, the temptation of patient agents to violate trust can be kept in check by scaling down initial cooperation to

$$x_0 = \frac{\delta}{[1 - \delta(1 - \pi)][\pi\beta + (1 - \pi)\gamma] + \delta[\pi - (1 - \pi)\gamma]},$$

or to “almost zero” using our convention if the initial experiment is unprofitable for patient types. For derivations, see the Appendix. The intuition is as follows. If (11) is violated, the temptation to cheat is too strong as $x_0 \rightarrow 1$. At the opposite end, as $x_0 \rightarrow 0$, this temptation is reversed: the stakes are too small for cheating strangers to generate much profit. There is always some intermediate value of x_0 where the no-deviation constraint binds, and this value gives us the initial cooperation level under an efficient norm. It also satisfies bilateral rationality, since any attempt to renegotiate to a higher stake at the outset will be met with individual deviations.

5. BILATERAL RATIONALITY WITH MANY TYPES

The two-type model, with one type entirely myopic, is obviously highly stylized. Agents may have a wide range of privately known discount factors, and it is reasonable to suppose that this

⁹This is necessary to get cooperation in the bilateral repeated game. Assuming perpetual defection as punishment, the no-deviation constraint in the repeated game is $(1 - \delta)\beta \leq \delta$, which on rearrangement gives us the bound on δ .

would lead to more gradual learning (in several steps, not just two) about a partner's traits. But the interplay between private arrangements at the level of the matched pair and social aspects of the norm such as adequate prevention of non-compliance becomes a more nuanced and complex issue.

This isn't a question of generalization for its own sake. Certainly, such an investigation serves as a robustness check on the qualitative results of the two-type case, but it will also yield additional insights. One might ask, for instance, if cooperative paths invariably settle down to a limit, and the extent of discount-mixing that can remain among matched pairs in that limit. After all, the welfare properties of a sustainable norm will depend on the extent to which it *eventually* enables players to match assortatively — perhaps after trying out many relationships. One can study the extent of cooperation (after the gradualist phase) as the distribution of discount factors moves towards greater or lesser patience. And certainly, it is possible to combine gradualism with other more traditional means of sanction, such as the observation and dissemination of past conduct via record-keeping. By parametrically changing the cost of such record-keeping, it is possible to consider not just the two extremes of complete information and no information, but to understand how society morphs from one arrangement to another as the cost of record-keeping changes.

However, as we've already noted, extending the concepts of efficiency, incentive compatibility and bilateral rationality to such a set-up is not straightforward. Below, we sketch the outline of a framework, highlighting some conceptual difficulties that lie along the way and proposing what we think are reasonable analytical choices that could be made. This is a topic of our ongoing research and our aim here is to suggest a way forward rather than present any definitive results.

5.1. Sustainability. Recall from Section 4.2 that a norm \mathbf{x} is more than a mere specification of a path: it prescribes $x_t(\boldsymbol{\pi})$ and an exit decision $d_t(\boldsymbol{\pi}, \delta) \in \{0, 1\}$ for each stage t , belief $\boldsymbol{\pi}$ and type δ . To prepare for the formulation of individual and bilateral rationality with many types, we explicitly note that a matched pair of agents will have a plan \mathbf{y} of its own, not necessarily adhering to the ambient norm \mathbf{x} . Mathematically, \mathbf{y} looks just like a norm: it prescribes cooperation levels $y(\boldsymbol{\pi})$ for every belief at the start of any stage, and individual prescriptions to default or exit $e(\boldsymbol{\pi}, \delta) \in \{0, 1\}$, that depend both on privately known type as well as the public belief summarized by $\boldsymbol{\pi}$. In the sequel, the x - and y -functions, and the d - and e -functions, must coincide. That is, we ask if any matched pair will willingly follow \mathbf{x} .

As before, we first define continuation values. An individual who is in a matched relationship and follows the norm \mathbf{x} will have a sequence of continuation values $V_t^{\mathbf{x}}(\boldsymbol{\pi}, \delta)$ at any stage t , assuming that the current belief is $\boldsymbol{\pi}$ and assuming all exit recommendations are followed. Next, an individual who is in a matched relationship and follows the plan \mathbf{y} (but with ambient norm \mathbf{x}) will have a sequence of continuation values $V_t^{\mathbf{x}\mathbf{y}}(\boldsymbol{\pi}, \delta)$ at any stage t , where it is understood that exiting the current relationship will give rise to a continuation $V_0^{\mathbf{x}}(\boldsymbol{\pi}, \delta)$ in the period after. The incentive-compatibility of the exit recommendations can now be checked. *For agents currently following \mathbf{x} :* for every δ and for every t , if:

$$(1 - \pi_t) \underbrace{[(1 - \delta)\ell(x_t) + \delta V_0^{\mathbf{x}}(\boldsymbol{\pi}_0, \delta)]}_{\text{cooperate, partner exits}} + \pi_t \underbrace{[(1 - \delta)v(x_t) + \delta V_{t+1}^{\mathbf{x}}(\boldsymbol{\pi}', \delta)]}_{\text{both cooperate}} \geq \underbrace{(1 - \delta)\pi_t[v(x_t) + g(x_t)] + \delta V_0^{\mathbf{x}}(\boldsymbol{\pi}_0, \delta)}_{\text{exit}},$$

then $d_t(\boldsymbol{\pi}, \delta) = 0$, and if the opposite strict inequality holds, then $d_t(\boldsymbol{\pi}, \delta) = 1$,¹⁰ where π_t is the conditional probability of partner cooperation given the recommendations under \boldsymbol{x} for that stage and the going prior $\boldsymbol{\pi}$, and $\boldsymbol{\pi}'$ is the posterior on the partner conditional on no exit. The left-hand side is the payoff to type δ if she continues at stage t , and it includes the possibility that she might herself be “cheated upon.” The right-hand side is similarly the expected payoff from exit. If the former exceeds the latter, the condition dictates that our individual continues the relationship at that date; otherwise she exits. In particular, the condition incorporates the one-shot deviation principle so familiar from dynamic programming, because future exit recommendations (that determine continuation values) are assumed to be honored. We can rewrite this condition to obtain a sustainability restriction akin to that for complete information: $d_t(\boldsymbol{\pi}, \delta) = 0$ if and only if

$$(12) \quad \delta[V_{t+1}^{\boldsymbol{x}}(\boldsymbol{\pi}', \delta) - V_0^{\boldsymbol{x}}(\boldsymbol{\pi}_0, \delta)] + \frac{1 - \pi_t}{\pi_t}(1 - \delta)\ell(x(t)) \geq (1 - \delta)g(x(t)),$$

where the additional second term on the left-hand side captures the possibility of being “suckered” with a negative payoff in case of cooperative continuation.

For agents currently following \boldsymbol{y} : all we need to is adjust (12) at just one place, so that the continuation value on the left hand conditional on continuation is $V_{t+1}^{\boldsymbol{x}\boldsymbol{y}}$ instead of $V_{t+1}^{\boldsymbol{x}}$. Incentive-compatibility then requires that $e_t(\boldsymbol{\pi}, \delta) = 0$ if and only if

$$(13) \quad \delta[V_{t+1}^{\boldsymbol{x}\boldsymbol{y}}(\boldsymbol{\pi}', \delta) - V_0^{\boldsymbol{x}}(\boldsymbol{\pi}_0, \delta)] + \frac{1 - \pi_t}{\pi_t}(1 - \delta)\ell(x(t)) \geq (1 - \delta)g(x(t)),$$

where π_t is now to be interpreted as the partner’s probability of cooperation at stage t under \boldsymbol{y} .

Our focus on general equilibrium is specifically captured by the presence of the term $V_0^{\boldsymbol{x}}(\boldsymbol{\pi}_0, \delta)$, which represents the value to an individual from embarking on a new relationship. This is not some exogenously specified object. It is given by the norm itself, which is rooted both in society and in the willing responses of matched pairs. Notice also how under this formulation, different discount factor types will have different, endogenously varying outside options.

5.2. Bilateral Rationality. We now revisit the question of bilateral rationality. As already noted, in the ambience of some social norm \boldsymbol{x} satisfying (12), a matched pair will freely choose its own plan \boldsymbol{y} of cooperation and exit. In equilibrium, we ask for that freely chosen path to coincide with \boldsymbol{x} . Now say that \boldsymbol{x} is *vulnerable to bilateral renegotiation* if for some matched pair, stage t and (on-path) prior $\boldsymbol{\pi}$ there is an alternative plan \boldsymbol{y} , starting from stage t , with its own (incentive-compatible) exit strategy as described by (13), such that every type under $\boldsymbol{\pi}$ is better off adhering to \boldsymbol{y} rather than \boldsymbol{x} .¹¹ More succinctly, if for some t , (on-path) prior $\boldsymbol{\pi}$, and \boldsymbol{y} satisfying (13),

$$V_t^{\boldsymbol{x}\boldsymbol{y}}(\boldsymbol{\pi}, \delta) > V_t^{\boldsymbol{x}}(\boldsymbol{\pi}, \delta)$$

for all δ in the support of $\boldsymbol{\pi}$, then \boldsymbol{x} is vulnerable to bilateral renegotiation. For instance, if for some t , $x(t)$ can be safely increased without causing any additional exit at stage t — or in other words, if (12) is slack at stage t for some going $\boldsymbol{\pi}$ — then that surely suffices for the vulnerability of \boldsymbol{x} to bilateral renegotiation. We will say that a sustainable norm is *bilaterally rational* if it is not vulnerable to bilateral renegotiation at any on-path belief and at any stage.

¹⁰We presume continuation under indifference. This is a property typically exhibited in equilibrium, anyway, even if we do not presume it.

¹¹This is akin to internal consistency, introduced in Bernheim and Ray (1989) and Farrell and Maskin (1989).

We make three remarks. First, in these contemplated deviations, the social norm x itself stays unchanged. So for instance, when $x(t)$ is raised in the last line of the previous paragraph, it is only raised in the relationship — in the plan y as it departs from the norm — and not in the norm, which is unaltered by this “bilateral deviation.” One pair cannot affect society. Second, so as to accommodate various maximands, the renegotiation concept is very demanding (“all types must be better off”) which means that the restriction imposed by it is weak. Third, in the definition we only examine vulnerability for on-path beliefs; namely, beliefs that are reached by applying the norm and its exit recommendations at each stage, starting from the prior π_0 . One could apply the definition to all beliefs, reached on-path or not, and that would increase the scope for vulnerability, but would not have substantively different implications.

The reader might wonder why such a constraint needs to be imposed at all. Wouldn’t Pareto-optimality across partners at date 0 imply Pareto-optimality at future dates? The answer is generally in the negative. A partnership might *want* to keep $x(t)$ low (for some future t) in order to induce swift exit today by depressing continuation values for some unwanted impatient types. Such a strategy would require a commitment on the part of the partnership to keep $x(t)$ low when stage t rolls around. But this could be impossibly hard to do; at any rate, that is our presumption. If, for instance, for every remaining δ -type at t , the enforcement constraint (12) is slack, the urge to renegotiate $x(t)$ upwards, *taking x as given for the rest of society*, could be irresistible. We therefore impose renegotiation-proofness as a constraint on the “best response” of the matched pair to the going norm, whatever specific form that best response might take.

5.3. Brief Remarks on Alternative Foundations for Bilateral Rationality. Recall that in the two-type case with one type “completely impatient” ($\delta = 0$), we settled on an efficiency concept that maximized the expected payoff to the patient type. The underlying rationale for this can be expressed by the following informal speech, which either party could make to the other:

“The impatient type is someone we’d like to weed out, so why care about that type’s welfare in the first place? Let’s just extend a small amount of cooperation to begin with (to make sure we’re not that type), and then we can proceed to a cooperative steady state that’s best for our patient selves. Please don’t tempt fate by asking me to place any welfare weight on the impatient type: you know exactly what sort of suspicions *that* will arouse. The only reasonable thing for us to do is to try and solve for the best path that we would want, *were we both to be the patient type.*”

There is something focal about the patient type when there are just two types and the other type is an inveterate deviator. It is natural to think that both players would want to approach the relationship as if they are the patient type. Anything else would come across as suspicious, to say the least. So the choice of maximizing the payoffs to the patient type is possibly the only arrangement which makes sense under the umbrella of bilateral rationality. But this focus is immediately lost when there are many types. Imagine, for instance that the distribution of types have full support on $[0, 1]$. Would it make sense to emphasize the infinitely patient types at the expense of everyone else? We have therefore retreated to the more eclectic device of the Pareto criterion.

But other options are possible. An obvious alternative is for society to possess a common weighting scheme $\{\omega(\delta)\}$ defined on the payoffs of all types. Those weights could be increasing in δ if patience is a valued trait, and it is additionally possible that $\omega(0) = 0$, just as in the two-type

model. Certainly, the specific form of a sustainable and bilaterally rational norm will then depend on the weighting structure, but to the extent that its qualitative features are robust to the specific choice of weights, this is a potentially useful alternative line of inquiry.

A different option that also nests the two-type case as exposited here is an argument based on inclusion and exclusion of types at any stage. Specifically, imagine that a matched pair has a current belief π over their types. Their choice of cooperation level will cause some subset of their type space to defect, and the remaining subset to cooperate. Then an extension of the speech above could read:

“This is the subset of types that we’d like to weed out at this stage, so why care about the welfare of any of them? We fully expect them to defect, after all. Let us maximize the sum of expected payoffs to the cooperative types instead.”

The same argument would be repeated at future stages, as long as there are types to be deliberately excluded. We do not pursue this option here.

5.4. The Structure of Bilaterally Rational Norms. While bilateral rationality can take many forms — and we have chosen a particularly weak version — it goes a long way towards disciplining the structure of norms.

The first feature is the implied stationarity of any sustainable, bilaterally norm when beliefs are unchanging. The intuition behind this is very similar — though not identical — to the logic that challenged gradualist norms in the complete-information case. There, a non-stationary norm was undermined by its very non-stationarity, as matched pairs would want to renegotiate to the most advantageous *stationary* cooperative level, given that norm. The same is true with incomplete information as long as the support over δ -types is unchanging across stages. After all, a higher value of x is beneficial to all such types. Therefore a pair faced with an unchanged on-path belief along the norm would immediately wish to negotiate to the best cooperative level supported by that norm, installing that level as a stationary plan y .¹² In equilibrium, $y = x$, so this is true of x as well.

The second feature, a near-corollary of the first, is that all sustainable and bilaterally rational norms must eventually converge to some well-defined limit level of cooperation. For the set of on-path beliefs is forever shrinking along any norm, and so must converge to some limit set. Far enough out on that sequence, then, the path itself cannot exhibit a degree of fluctuation that is bounded away from zero. The steadily dampening fluctuation proves convergence to a limit. We think of this as the analogue to the stationary cooperative path among patient types in the two-type case.

The remaining features require more subtle analysis, and we merely summarize our ongoing research in the form of looser conjectures.

Third, every sustainable and bilaterally rational norm must converge *monotonically upwards* to the limit identified above. For if $x_t > x_{t+1}$ at any stage, then it must mean that continuation values along the norm just following x_t must strictly exceed continuation values just following x_{t+1} . (For

¹²Because the set of types is unchanging, and the social norm is presumed to be sustainable, all types would prefer to continue along the norm when the current cooperation level — and attendant temptation to deviate — is highest, or more precisely, close to its supremum value. But then they must *a fortiori* wish to continue when all future cooperation levels are replaced by that supremum, which proves that non-stationarity is vulnerable to bilateral renegotiation.

if this were false, then the incentive constraint at stage $t+1$ must be slack for all types that complied with the norm at date t , opening the door to upward renegotiation of the norm at stage $t+1$.) But now this downward drift in continuation values can be replaced by a new plan starting at date $t+1$ which restarts the earlier norm at date t , thereby raising continuation values relative to those prescribed by the norm at $t+1$.¹³

Fourth, with monotonicity established, it is possible to show that such monotonicity is *strict* until the first point at which the norm becomes stationary, upon which it is stationary for ever after. Moreover, that strict phase must be nonempty for any sustainable and bilaterally rationally norm which involves non-trivial cooperation. The first of these observations is established by showing that any “temporary stationarity” along the path is vulnerable to bilateral renegotiation: the more desirable rising phase can simply be brought forward in time, raising continuation values. The second observation follows simply from the fact that nontrivial levels of cooperation can only be sustainable if deviation implies lower continuation value under the restarted norm, which cannot be the case if the norm is perennially stationary.

Taken together, the third and fourth features bring out the essence of gradualism, which is that a norm exhibits steadily increasing levels of cooperation over time before leveling out.

Our final points are relatively nebulous at this stage, but worth stating here as useful directions of research. As motivation, observe that once the norm prescribes stationary cooperation levels, the set of cooperative δ -types becomes stationary and unchanging as well. (This is true in an asymptotic sense in case the norm converges “at infinity.”) One might ask what this set looks like, and what that means for the social efficacy of matching. The smaller is this final set, the more finely — or more “assortatively” — matched is the set of final compliant types, but is that necessarily a good thing? For this must also mean that the remaining types remain in a perennial cycle over the norm, periodically splitting off from their relationships and randomly matching with one another. There is, therefore, a tradeoff between the extent of gradualism in the norm and its ability to bring closely matched types together, and the amount of “churning” generated for all other types. This can only be evaluated by evaluating these outcomes in the context of the overall distribution of types. If that distribution places high probability on the limit set of compliant types, then we have a well-functioning norm; otherwise, a coarser, less ambitious norm may well do better.

We end by asking the seemingly most obvious of these questions: are the most patient types the ones most likely to cooperate along a sustainable, bilaterally rational norm? It is certainly true that patient types are more likely to value the future levels of cooperation provided by such a norm. And yet, they are also the ones who are least deterred by the prospect of gradualism, as they are future-oriented. The punishment value of restarting is therefore lower. This apparent ambiguity can indeed be shown to be resolved in favor of the patient types, so that the limit set is of the form $[\delta^*, 1]$ for some $\delta^* \in (0, 1)$. But this result is non-trivial and in fact, it is generally false for norms that are not monotone in cooperation levels (though the third feature of sustainable, bilaterally rational norms described above will eliminate this possibility).

With this final result in hand, we can fully characterize the limit level of cooperation, which then permits a variety of explorations as various parameters of the model are changed. The limit level

¹³The argument is more subtle than this because it requires us to show that continuation values move in the same direction for all types along a norm, something that is obviously true when *stationary* cooperation levels are adjusted up or down, as in the first feature discussed earlier.

of cooperation — call it x^* — must be the one that makes the threshold δ^* -type indifferent between ongoing cooperation and defection into the unmatched pool. For all other types $\delta > \delta^*$, the incentive constraint is slack. If types are continuously distributed, that raises the question of whether a little more experimentation — a little increase in x^* — is worth it. The potential gains come from a still higher level of cooperation thereafter. The potential losses come from the heightened possibility that the norm will be broken by a set of types located very close to δ^* . It is this balancing of gains and losses at the margin that will jointly pin down the limit (x^*, δ^*) , an exercise that will depend in a rich way on all the parameters of the model.

5.5. Other Dimensions of Heterogeneity. Reputation effects in our model arise from heterogeneous discount factors. One could imagine agents differing on other aspects of payoffs — the gain $g(x)$ from cheating or the loss $l(x)$ from being cheated, for example. The main advantage of having a type with $\delta = 0$ is that playing the action 0 is a dominant strategy for such a type. That simplicity is lost if we consider differences on these other dimensions. Nevertheless, given that equilibria with heterogeneous norms exist even in the one-type model (see Section 7.1 below), it should be possible to construct similar equilibria here, with agents having higher gain or loss adopting the role of serial cheats. An added benefit is that each type will strictly prefer its own strategy to the other's, instead of being indifferent. Note that the bad types (the ones facing higher gain/loss) will have no unilateral or bilateral incentive to try and form long term cooperative bonds with their own types if the initial experimental level of cooperation offered by the good types is lucrative enough to exploit. As before, parameter restrictions are needed for existence.

In the two-type models considered so far, cooperation jumps to a steady-state immediately after clearing stage 0 successfully. Macchiavelo and Morjaria (2015) consider a variant where there is more gradual cooperation build-up even with two types. In their model, the bad type has the same preferences as the good type, but is compelled to play the dominant action 0 in some periods for exogenous reasons. In this environment, continued reciprocation leads to slow Bayesian updating, permitting more gradualist paths.

6. COMPLETE-INFORMATION ROUTES TO ANONYMOUS COOPERATION

We end this paper with some remarks on alternative approaches to building cooperation in large anonymous societies. These approaches do not rely on incomplete information but on other aspects of the complete information environment.

6.1. Heterogenous Norms. One way to mimic a mix of patient and impatient agents without resorting to incomplete information is to have identical agents behave in non-identical ways. Specifically, even when all agents have the same discount factor (so there is complete information), the emergence of cooperative norms is possible if we allow them to adopt different norms or strategies. An equilibrium is reached at the population level if the mixture of strategies is such that *all the strategies employed yield the same expected payoff*, as in polymorphic ESS in the evolutionary biology literature. This is the approach taken in Okuno-Fujiwara and Fujiwara-Greve (2009).

To illustrate the idea as simply as possible, consider the linear example in Section 4.4 with binary actions, i.e., $x \in \{0, 1\}$. Consider two strategies: *trustworthy* and *untrustworthy*. The trustworthy strategy involves always playing $x = 1$, and forming or maintaining a relationship as long as cooperative behavior is reciprocated. The untrustworthy strategy involves always choosing $x = 0$

and never continuing a relationship. The lifetime expected payoffs from these strategies (at the beginning of a new relationship) depend on π , the fraction of agents who choose to adopt the trustworthy strategy, and in the example are given by

$$\begin{aligned} V_0^u &= (1 + \beta)\pi \\ V_0^t &= \frac{\pi - \gamma(1 - \delta)(1 - \pi)}{1 - \delta(1 - \pi)} \end{aligned}$$

While the payoffs of untrustworthy agents are linear in π , that of trustworthy agents is concave. Under suitable parametric restrictions, there are two points of intersection, both of which are candidates for a mixed population equilibrium.

One reasonable adjustment dynamic is the following: π grows if $V_0^t > V_0^u$, and decreases if $V_0^t < V_0^u$. Under this dynamic, there are two locally stable population compositions: $\pi = 0$ and the higher of the two solutions to $V_0^t = V_0^u$, which is given by

$$\pi = \frac{1 - (1 - \delta)(1 + \beta - \gamma) + \sqrt{[1 - (1 - \delta)(1 + \beta - \gamma)]^2 - 4(1 + \beta)\gamma\delta(1 - \delta)}}{2\delta(1 + \beta)}$$

Thus, behavioral heterogeneity may emerge endogenously and successful cooperative norms succeed only when there is less than universal adherence to them. Extending this kind of “evolutionary” analysis to the more general model could be a topic for further research. The interesting aspect of this generalization is the non-trivial interaction between gradualism (the path \mathbf{x}) and the endogenously determined proportion of cooperators (π).

We remark that this sort of polymorphic approach sits uncomfortably with the notion of bilateral rationality that we have explored in the main analysis of the paper. It exploits the indifference of agents to have some of them deviate in the relationship. A bilaterally rational pair can avoid this outcome by choosing an arrangement that retains a “small” surplus to each partner, so as to make it strictly worthwhile for each of them to continue the relationship. Technically, we would model this as a lexicographic preference for cooperation when an agent is indifferent — after all, non-compliance creates no extra gain for the deviating agent in the face of her indifference, but it does create a first-order loss for her matched partner.

6.2. One-Sided Moral Hazard and Transferable Utility. In many applications, moral hazard is one-sided — recall the asymmetric principal-agent problem described earlier in Section 2.2. Furthermore, monetary payments can be made from the principal to the agent — or vice versa — can complement gradualism in curbing opportunistic behavior. Examples include employment relationships (where the worker can shirk), credit relationships (where the borrower may default), and transactions in goods (where the seller may undersupply quality). Wages, interest rates or product pricing play a critical role in relational arrangements in these scenarios.

One idea explored in the literature is that of a price premium or efficiency wage. The agent who can indulge in morally hazardous behavior is given a monetary transfer over and above what her bargaining power would have extracted were contracts fully enforceable. The fear of losing this premium may be enough to prevent opportunism when an agent’s past behavior is publicly known (Klein and Leffler 1981). However, in anonymous environments, it is obvious that a price premium, when paid upfront as an efficiency wage, or simply promised such as an interest rate discount, cannot generate the right incentives. If agents could form a new relationship without cost or delay, they will keep enjoying the premium in new relationships, even after the termination of an ongoing

relationship, because past malfeasance is not detected. However, in many models, price premia can combine with rationing or matching frictions to produce the desired disciplining effect. In some models, the matching friction is exogenous (Greif 1993); in others (such as Shapiro and Stiglitz 1984 or Eswaran and Kotwal 1985), it arises endogenously from equilibrium interaction.¹⁴

In most such models, the choice of cooperation level is modeled as binary. The worker is either hired or not; the borrower is given a loan (of fixed size) or denied credit; the consumer either buys the indivisible good or refrains from buying. That largely precludes gradualism as a disciplining device. Datta (1995) and Wei (2018) do study one-sided moral hazard problems with variable cooperation levels but their models do not have transfers. In Ghosh and Ray (2016), we allowed for all three instruments in a model of the informal credit market — variable interest rates (or premia), variable loan size (or cooperation levels) and the possibility of credit rationing (or endogenous rematching frictions). Some borrowers are modeled as habitual defaulters, but the credit histories of new clients are not available to lenders, so the central forces of the two-sided moral hazard model are also present there. A combination of price premia and gradualism (i.e., lower interest rates and relaxed credit limits in the mature phase of a credit relationship) arise in equilibrium as disciplinary devices, sometimes complemented by the rationing of borrowers.

Ray (2002) highlights another way in which *non-stationary* monetary transfers can play an incentivizing role. This involves allowing the agent to keep all the surplus arising from the relationship towards its later stages, while the principal extracts steep payments upfront. This backloading of rewards is very much in the spirit of gradualism in our setting, except that it is brought about through a zero sum instrument. More crucially, Ray assumes that the agent receives an exogenous outside option upon break-up of the relationship, so there are no general equilibrium effects at all. Given that our focus here is on the society-wide ramifications of anonymous relational contracting, Ray's model must be expanded so that the outside option is endogenous at the societal level.

Consider, then, matched pairs of principal and agent with common $\delta \in (0, 1)$. Suppose that at each stage, the principal *commits* to both the scale of the interaction at some $x \in [0, \bar{x}]$, and a nonnegative transfer τ (there is agent limited liability). If the agent acts honestly, the principal gets a net payoff of $p(x) - \tau$, where $p(x) > 0$ for $x > 0$, while the agent gets $a(x) + \tau$, where $a(x)$ may well be negative; for instance, if she expends effort. If the agent acts dishonestly, she gains $g(x) + (\mu - 1)\tau$, while the principal obtains a payoff of $l(x) - \mu\tau$, where $\mu \in [0, 1]$ is the fraction of the transfer that the agent can retain upon deviation. This accommodates settings in which the transfer τ is made upfront, in which case $\mu = 1$ and the gain is independent of τ , or the transfer is made only conditional on performance, in which case $\mu = 0$ and the gain must net out the foregone transfer. The principal's outside option is normalized to zero. We maintain the focus of this paper by assuming that there are no matching frictions, and there is full anonymity of agent histories.

Recall that if contracts were enforceable by third parties, the scale of interaction every period would be at the first-best level, $x^1 = \arg \max_x [p(x) + a(x)]$, regardless of the distribution of bargaining power. With non-contractible compliance, a self-enforcing contract is some path of cooperation

¹⁴Endogenous frictions arise when the price premium affects the number of relationships that principals want to create. For example, efficiency wages (wages above the competitive level) may not only discipline workers but also reduce labor demand and give rise to equilibrium unemployment, as in Shapiro and Stiglitz (1984) or Macleod and Mslcomson (1989). Other papers which discuss solutions to the problem of shirking include Macleod and Malcomson (1998), Yang (2008) and Fahn and Murooka (2022).

and payments, $\{x_t, \tau_t\}_{t=0}^{\infty}$, satisfying the agent's no-deviation constraints for all t :

$$(1 - \delta)[g(x_t) - (\mu - 1)\tau_t] \leq \delta [V_{t+1} - V_0]$$

where

$$V_t = \sum_{s=t}^{\infty} \delta^{s-t} [a(x_s) + \tau_s]$$

is the stream of future net payoffs expected by the agent if she remains compliant.

If the agent had all the bargaining power, she would demand and receive a payment of $\tau = p(x^1)$ every period, and so capture the entire surplus. With self-enforcement constraints, the first-best is generally not achievable. With assumptions on p and a , it is easy to see that in this case, the level of cooperation would be stationary at some level $x^* < x^1$, with transfers chosen to meet the principal's participation constraint, which is taken to mean non-negative payoff.

Now move bargaining power from agent to principal. Then, with limited liability on the part of the agent, the contract will generally start small and build over time. If the principal has enough power so that the agent's limited liability constraint bites, the agent will in effect become an apprentice, working for free during some initial stages of the relationship. Meanwhile, it can be shown (Ray 2002) that the scale of the relationship pushes upward, and that *irrespective of the assignment of bargaining power*, must eventually attain the level x^* , which is the agent's constrained optimum when *she* has all the bargaining power. Indeed, all relational arrangements exhibit gradualism except the one that pertains to agent-maximal bargaining power, but they all have the same limit x^* , which corresponds to the stationary arrangement with agent-maximal power. In short, contractual "tails" are invariant to the distribution of power.

But the extent of gradualism is inversely related to the agent's power. As a special case, consider the norm that results when the principal has all the power. We are now looking at the opposite corner of the constrained Pareto frontier, where the agent's payoff is now at a minimum relative to all points on that frontier. Nevertheless, as Ray (2002) shows, the resulting principal-optimal path must, after some date, maximize the agent's payoff, with $x_t = x^*$ and $\tau_t = p(x^*)$ from that date onward. The principal makes all her money before that date. If the agent had no limited liability, this could be done "selling the firm" upfront to the agent.

Consider the general equilibrium implications of these observations. In a standard setting with fixed outside option for the agent, the story ends as described above. When those outside options are generated by the agent's re-entry into a pool with subsequent matching, her value V_0 will be negatively related to the principal's power. Relational contracts in this environment will therefore generate greater aggregate surplus when principals have greater bargaining power. In particular, x^* will be pushed *upwards* in all paired relationships when principals have more power, an effect that is entirely a product of the general-equilibrium interactions in the model.

This finding stands in contrast to static principal-agent models with limited liability when bargaining power is exogenously altered. Mookherjee (1997) observed that with greater agent power, the optimal contract generates greater incentives to the agent, thereby increasing the scale and efficiency of operations. Ghosh, Mookherjee and Ray (2000) find similar results in the context of credit market relationships. In dynamic models with random matching across principal and agent in every period, the same observation is true. With greater agent power, there are salutary dynamic

effects via the agent's incentives to save, allowing her to possibly escape poverty traps (Mookherjee and Ray 2002). In all these settings, greater power to the agent enhances overall efficiency.

On the other hand, in the current framework, principals obtain their pound of flesh via backloading contractual payments. Stronger principals are able to postpone the sharing of surplus with the agent to a more distant future and will prolong rent extraction at the beginning of the relationship. This lowers the agent's payoff whenever she starts a new relationship, and consequently worsens her outside option — falling out with a current principal implies renewing a long period of toil with another strong principal. The resultant slackening of the incentive constraint *at every stage of the relationship* permits a move towards a higher scale of productive activity without the fear of the principal getting suckered. Now greater power to the *principal* enhances overall efficiency. These opposing effects are of interest, and a deeper exploration of the connection between bargaining power and overall efficiency must await a more detailed investigation.

7. CONCLUSION

We have studied the dynamics of individual relationships when cooperation is driven by the prospect of entering back into the same class of relationships, as opposed to some explicit threat of punishment. This formulation, based on a self-referential equilibrium path, is particularly apt for situations in which a history of past behavior (ranging from economic misdemeanors such as bad business practices or loan defaults to social ills such as domestic abuse) *is not observable*, or is only observable at high cost. As the title of our paper indicates, our emphasis throughout has been on the wider “general equilibrium” consequences of the resulting relational arrangements.

We believe that this sort of investigation contributes in a serious way to the theory of informal institutions, a phrase that is often bandied about with little real content. We often think of formal institutions as social arrangements that achieve — or at least are designed for — second-best efficiency in a world plagued by problems of adverse selection and/or moral hazard. That design-based approach makes sense when an institution is brought into being by a deliberate policy, which is ratified — with its social ramifications explicitly in mind — by a vote in Congress or a Court decree. But what of *informal* institutions, that emerge not from parliamentary or judicial deliberations, but spontaneously and in a decentralized way from everyday actions? What is the underlying teleology of such institutions — why would they necessarily gravitate in some directed way toward constrained social efficiency?

There are two broad approaches to this fundamental question. One is that the efficiency of informal institutions is shaped by evolutionary encounters across societies — that societies compete as evolutionary units and develop good institutions as a result of high-level natural selection. The other is to study the decentralized features of such institutions and ask if those features — perhaps rational or rationalizable at the level of individuals or small groups — possess serendipitous externalities that are conducive to social efficiency. Our investigation into the “general equilibrium of relational arrangements” is precisely in the spirit of this second approach. From a social perspective, a norm of behavior must provide minimally adequate incentives for cooperation in relationships, economic or otherwise. From the private perspective of any one relationship, the norm of behavior elsewhere is taken as given, while interactions within the relationship are designed for the well-being of specific participants. When might the latter “micro” aspect of individual convenience be in conformity

with, or antagonistic to, the former “macro” imperative of social suitability? That is the kind of question that a study of the general equilibrium of relational arrangements seeks to answer.

REFERENCES

- Abreu, D., Pearce, D., and Stacchetti, E. (1990): “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring.” *Econometrica*, 1041-1063.
- Aleem, I. (1990): “Imperfect Information, Screening, and the Costs of Informal Lending: A Study of a Rural Credit Market in Pakistan,” *World Bank Economic Review*, 4, 329-349.
- Bendor, J., and Mookherjee, D. (1990): “Norms, Third-party Sanctions, and Cooperation.” *Journal of Law, Economics, and Organization*, 6(1), 33-63.
- Bardhan, P. (Ed.). (1989): *The Economic Theory of Agrarian Institutions*. Clarendon Press.
- Bernheim, D. and D. Ray (1989), “Collective Dynamic Consistency in Repeated Games,” *Games and Economic Behavior*, 1, 295–326.
- Besley, T., Coate, S., and Loury, G. (1993): “The Economics of Rotating Savings and Credit Associations.” *American Economic Review*, 83(4), 792-810.
- Bhaskar, V., and Thomas, C. (2019): “Community Enforcement of Trust with Bounded Memory.” *Review of Economic Studies*, 86(3), 1010-1032.
- Bodoh-Creed, A. L. (2019): “Endogenous Institutional Selection, Building Trust, and Economic Growth.” *Games and Economic Behavior*, 114, 169-176.
- Carmichael, H. L. and MacLeod, W. B. (1997): “Gift Giving and the Evolution of Cooperation.” *International Economic Review*, 485-509.
- Clark, D., Fudenberg, D., and Wolitzky, A. (2021): “Record-Keeping and Cooperation in Large Societies.” *Review of Economic Studies*, 88(5), 2179-2209.
- Coate, S. and Ravallion, M. (1993): “Reciprocity without Commitment: Characterization and Performance of Informal Insurance Arrangements.” *Journal of Development Economics*, 40(1), 1-24.
- Datta, S. (1996): “Building Trust.” No. 1996/305, Suntory and Toyota International Centres for Economics and Related Disciplines, LSE.
- Deb, J. (2020): “Cooperation and Community Responsibility.” *Journal of Political Economy*, 128(5), 1976-2009.
- Deb, J., Sugaya, T., and Wolitzky, A. (2020): “The Folk Theorem in Repeated Games with Anonymous Random Matching.” *Econometrica*, 88(3), 917-964.
- Eaton, J. and Gersovitz, M. (1981): “Debt with Potential Repudiation: Theoretical and Empirical Analysis.” *Review of Economic Studies*, 48(2), 289-309.
- Ellison, G. (1994): “Cooperation in the Prisoner’s Dilemma with Anonymous Random Matching.” *Review of Economic Studies*, 61(3), 567-588.

- Eeckhout, J. (2006): “ Minorities and Endogenous Segregation.” *Review of Economic Studies*, 73(1), 31-53.
- Eswaran, M., and Kotwal, A. (1985): “ A Theory of Two-Tier Labor Markets in Agrarian Economies.” *American Economic Review*, 75(1), 162-177.
- Fafchamps, M. and Minten, B. (1999): “Relationships and Traders in Madagascar.” *Journal of Development Studies*, 35(6), 1-35.
- Fahn, M. and Murooka, T. (2022): “Informal Incentives and Labor Markets.” Mimeo.
- Farrell, J. and E. Maskin (1989), “Renegotiation in Repeated Games,” *Games and Economic Behavior*, 1, 327–360.
- Fehr, E. and Schurtenberger, I. (2018): “Normative Foundations of Human Cooperation.” *Nature Human Behaviour*, 2(7), 458-468.
- Fieler, A. (2007): “ Repeated Partnership with Limited Information Flows.” Mimeo, University of Pennsylvania, https://drive.google.com/file/d/1H40OdWVbuqu3T3TI0udhHg_53YbKi3iD/view
- Frank, R. H. (1988): *Passions within Reason: The Strategic Role of the Emotions*. WW Norton and Co.
- Fujiwara-Greve, T. and Okuno-Fujiwara, M. (2009): “Voluntarily Separable Repeated Prisoner’s Dilemma,” *Review of Economic Studies*, 76(3), 993-1021.
- Ghosh, P. and Ray, D. (1996): “ Cooperation in Community Interaction without Information Flows.” *Review of Economic Studies*, 63(3), 491-519.
- Ghosh, P. and Ray, D. (2016): “ Information and Enforcement in Informal Credit Markets.” *Economica*, 83(329), 59-90.
- Ghosh, P., Mookherjee, D. and Ray, D. (2000): “ Credit Rationing in Developing Countries: An Overview of the Theory.” in Ray, D. and Mookherjee, D. (eds), *Readings in the Theory of Economic Development*, Blackwell, London.
- Greif, A. (1993): “ Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders’ Coalition.” *American Economic Review*, 525-548.
- Greif, A., Milgrom, P. and Weingast, B. R. (1994): “Coordination, Commitment, and Enforcement: The Case of the Merchant Guild.” *Journal of Political Economy*, 102(4), 745-776.
- Jackson, M. and A. Wolinsky (1996): “A Strategic Model of Social and Economic Networks,” *Journal of Economic Theory*, 71, 44-74.
- Johnson, S., McMillan, J. and Woodruff, C. (2002): “Courts and Relational Contracts.” *Journal of Law, Economics, and Organization*, 18(1), 221-277.
- Kandori, M. (1992): “ Social Norms and Community Enforcement.” *Review of Economic Studies*, 59(1), 63-80.
- Klein, B. and Leffler, K. B. (1981): “ The Role of Market Forces in Assuring Contractual Performance.” *Journal of Political Economy*, 89(4), 615-641.

- Kranton, R. E. (1996): "The Formation of Cooperative Relationships." *The Journal of Law, Economics, and Organization*, 12(1), 214-233.
- Kreps, D. M., Milgrom, P., Roberts, J. and Wilson, R. (1982): "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma." *Journal of Economic Theory*, 27(2), 245-252.
- Lindsey, J., Polak, B., and Zeckhauser, R. (2001): "Free Love, Fragile Fidelity, and Forgiveness: Rival Social Conventions under Hidden Information." Unpublished manuscript, Economics Department, Yale University.
- Macchiavello, R. and Morjaria, A. (2015): "The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports." *American Economic Review*, 105(9), 2911-2945.
- MacLeod, W. B. and Malcomson, J. M. (1989): "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica*, 57(2), 447-480.
- MacLeod, W. B. and Malcomson, J. M. (1998): "Motivation and Markets." *American Economic Review*, 88(3), 388-411.
- Mcauley, S. (1963): "Non-Contractual Relations in Business: A Preliminary Study." *American Sociological Review*, 28(1), 55-70.
- McMillan, J. and Woodruff, C. (1999): "Interfirm Relationships and Informal Credit in Vietnam." *Quarterly Journal of Economics*, 114(4), 1285-1320.
- Mookherjee, D. (1997): "Informational Rents and Property Rights in Land," in J. Roemer, ed., *Property Rights, Incentives and Welfare*. New York: Macmillan Press, 3-39.
- Mookherjee, D. and D. Ray (2002): "Contractual Structure and Wealth Accumulation," *American Economic Review*, 92, 818-849.
- Munshi, K. (2003): "Networks in the Modern Economy: Mexican Migrants in the US Labor Market." *Quarterly Journal of Economics*, 118(2), 549-599.
- Munshi, K., and Rosenzweig, M. (2009): "Why is Mobility in India so Low? Social Insurance, Inequality, and Growth." No. w14850, National Bureau of Economic Research.
- Okuno-Fujiwara, M. and Postlewaite, A. (1995): "Social Norms and Random Matching Games." *Games and Economic Behavior*, 9(1), 79-109.
- Ray, D. (2002): "The Time Structure of Self-Enforcing Agreements." *Econometrica*, 70(2), 547-82.
- Rosenthal, R. W. (1979): "Sequences of Games with Varying Opponents." *Econometrica*, 1353-1366.
- Rosenthal, R. W., and Landau, H. J. (1979): "A Game-Theoretic Analysis of Bargaining with Reputations." *Journal of Mathematical Psychology*, 20(3), 233-255.
- Ruffle, B. J., and Sosis, R. (2006): "Cooperation and the In-Group-Out-Group Bias: A Field Test on Israeli Kibbutz Members and City Residents." *Journal of Economic Behavior and Organization*, 60(2), 147-163.

Shapiro, C. (1983): "Premiums for High Quality Products as Returns to Reputations." *Quarterly Journal of Economics*, 98(4), 659-679.

Shapiro, C. and Stiglitz, J. E. (1984): "Equilibrium Unemployment as a Worker Discipline Device." *American Economic Review*, 74(3), 433-444.

Siamwalla, A., Pinthong, C., Poapongsakorn, N., Satsanguan, P., Nettayarak, P., Mingmaneeakin, W. and Tubpun, Y. (1990): "The Thai Rural Credit System: Public Subsidies, Private Information, and Segmented Markets." *The World Bank Economic Review*, 4(3), 271-295.

Udry, C. (1994): "Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria." *Review of Economic Studies*, 61(3), 495-526.

Watson, J. (1999): "Starting Small and Renegotiation." *Journal of Economic Theory*, 85(1), 52-90.

Watson, J. (2002): "Starting Small and Commitment." *Games and Economic Behavior*, 38(1), 176-199.

Watson, J. and Hua, X. (2022): "Starting Small and Renegotiation in Discrete-Time Relationships with a Continuum of Types." Mimeo.

Wei, D. (2019): "A Model of Trust Building with Anonymous Re-Matching." *Journal of Economic Behavior and Organization*, 158, 311-327.

Yang, H. (2008): "Efficiency Wages and Subjective Performance Pay." *Economic Inquiry*, 46(2), 179-196.

APPENDIX

LEMMA 1. Consider any sustainable pair (\mathbf{x}, d) with continuation values defined by $V_0 = V_0(\mathbf{x}) - d$ and $V_t = V_t(\mathbf{x})$ for all $t > 0$, as in the main text. Then $V_0 \leq V^*$, where V^* is defined in (4).

Proof. Consider any (\mathbf{x}, d) with associated values satisfying (2); then

$$(14) \quad (1 - \delta)g(x_t) \leq \delta[V_{t+1} - V_0] \text{ for all } t.$$

Define $\tilde{x} = \sup \{x_t\}_{t=0}^{\infty}$, consider any sequence of dates (with repetitions if the supremum is attained at some finite date) such that $x_t \rightarrow \tilde{x}$, and take a further subsequence so that V_{t+1} converges to some \tilde{V} . Obviously, $\tilde{V} \leq v(\tilde{x})$. Using all this information in (14) along with the continuity of g , we see that

$$(1 - \delta)g(\tilde{x}) \leq \delta[\tilde{V} - V_0] \leq \delta[v(\tilde{x}) - V_0],$$

so that

$$V_0 \leq \left[v(\tilde{x}) - \frac{1 - \delta}{\delta}g(\tilde{x}) \right] \leq V^*,$$

as claimed. ■

Proof of Proposition 1. By Lemma 1, $V_0 \leq V^*$. Additionally, it is easy to check that the construction given in the proposition generates $V_0 = V^*$. ■

Proof of Proposition 2. By the definition of T , $\delta^{T-1}(1 - \delta)v(0) + \delta^T v(x^*) \geq V^* > \delta^{T-1}(1 - \delta)v(x^*) + \delta^T v(x^*)$, so the existence of $\hat{x} \in [0, x^*)$ is guaranteed by the intermediate value theorem. By Lemma 1, we have $V_0 \leq V^*$ for all sustainable norms, so that $V_0 = V^*$ for an efficient norm given the construction in (7).

It is easy to check that the path constructed in (7) satisfies the no-deviation constraint (2) at every date (and makes it bind from date T onwards) and also achieves $V_0 = V^*$.

To show that any efficient norm \mathbf{x} must converge to some $x^* \in X^*$, define $\tilde{x} = \sup \{x_t\}_{t=0}^{\infty}$ as before. consider any sequence of dates (with repetitions if the supremum is attained at some finite date) such that $x_t \rightarrow \tilde{x}$, and take a further subsequence so that V_{t+1} converges to some \tilde{V} . As already noted, $\tilde{V} \leq v(\tilde{x})$. Passing to the limit in (2) and recalling that $V_0 = V^*$, we must have

$$(1 - \delta)g(\tilde{x}) \leq \delta[\tilde{V} - V^*]$$

which on rearrangement tells us that

$$(15) \quad V^* \leq \tilde{V} - \frac{1 - \delta}{\delta}g(\tilde{x}) \leq v(\tilde{x}) - \frac{1 - \delta}{\delta}g(\tilde{x}).$$

We claim that both inequalities in (15) must hold with equality, so that $\tilde{x} = x^*$ and $\tilde{V} = v(\tilde{x}) = v(x^*)$ for some $x^* \in X$. This follows immediately from comparing the first and last expressions in (15) and recalling (4). But the two equalities together imply that $\lim_t x_t$ exists and equals x^* . ■

Proof of Proposition 3. Index any norm by its lifetime return V_0 to the patient type at the start of a relationship, and consider the response of a matched pair which chooses its Stage-0 and later-stage

responses based on V_0 . At any Stage beyond 0, the matched pair will choose the largest stationary cooperation level y that meets (9) with equality; that is,

$$(16) \quad (1 - \delta)g(y) = \delta [v(y) - V_0].$$

with y_0 chosen to

$$(17) \quad \text{Maximize } \pi v(y_0) + (1 - \pi)l(y_0), \text{ subject to } (1 - \delta)[\pi g(y_0) - (1 - \pi)l(y_0)] \leq \delta [v(y) - V_0],$$

where y solves (16).¹⁵ These are well-defined if and only if $V_0 \leq V^*$; by the definition of V^* in (4), no y can satisfy (16) once $V_0 > V^*$. Define a mapping $\Psi(V_0, \pi)$ by

$$(18) \quad \Psi(V_0, \pi) = (1 - \delta)[\pi v(y_0) + (1 - \pi)l(y_0)] + \delta v(y),$$

where y and y_0 satisfy (16) and (17). For each π , the fixed points of Ψ in V_0 capture *all* bilaterally rational norms. But it is obvious that Ψ is declining in V_0 over the interval $[0, V^*]$: the value of y in (16) must shrink with V_0 and therefore so must the maximum value in (17). It follows that for each π a bilaterally rational norm (provided it exists) is unique up to the payoff V_0 it generates.

In fact, by the maximum theorem, which holds at every $V_0 < V^*$, Ψ declines *continuously* on $[0, V^*]$. Thereafter it is not defined as (16) can no longer be met. So a necessary and sufficient condition for existence (and associated uniqueness of equilibrium value V_0) is simply

$$(19) \quad \Psi(V^*, \pi) \leq V^*.$$

This condition is satisfied for an *interval* of the form $[0, \hat{\pi}]$. Clearly the “best response” y is invariant to π , given V_0 , while a decline in π must lower the maximum value in (17). Therefore, if (19) holds for some $\pi > 0$, it must hold for all lower $\pi' > 0$. To complete the proof, we note that $\hat{\pi} < 1$. For along any sequence π^k converging to 1, it is easy to see that $y_0^k \rightarrow y$ and therefore $\Psi(V^*, \pi^k) \rightarrow v(y)$. But by (16), we have $v(y) > V^*$, so (19) must fail for π sufficiently close to 1. ■

Linear Example: First, assume cooperation is binary, i.e., $x \in \{0, 1\}$. Consider the norm $\mathbf{x} = (1, 1, 1, \dots)$. Then $V_1 = 1$ and so, using (8),

$$V_0 = \frac{\pi - (1 - \pi)(1 - \delta)\gamma}{1 - \delta(1 - \pi)}.$$

The no-deviation constraint (10) for the patient type in Phase 0 can therefore be written as:

$$(1 - \delta) [\pi\beta + (1 - \pi)\gamma] \leq \delta [V_1 - V_0] = \frac{\delta(1 - \delta)(1 - \pi)(1 + \gamma)}{1 - \delta(1 - \pi)},$$

and manipulation of this inequality yields (11).

Now assume $x \in [0, 1]$. Consider a norm of the form $\mathbf{x} = (x_0, 1, 1, 1, \dots)$. Then $V_1 = 1$ and

$$V_0 = \frac{(1 - \delta)[\pi - (1 - \pi)\gamma]x_0 + \delta\pi}{1 - \delta(1 - \pi)}$$

In Phase 1 and beyond, the no-deviation constraint is satisfied by definition if $\delta \geq \underline{\delta}$. For it to be satisfied in Phase 0, we need

$$(1 - \delta) [\pi\beta + (1 - \pi)\gamma] x_0 \leq \delta [V_1 - V_0] = \frac{(1 - \delta) [1 - \{\pi - (1 - \pi)\gamma\}x_0]}{1 - \delta(1 - \pi)}.$$

¹⁵Recall our convention that $y_0 = 0$ if the maximand is negative for every $y_0 > 0$.

If initial experimentation has negative expected payoff; that is, if $\pi - (1 - \pi)\gamma < 0$, we can use our convention to think of x_0 as almost zero, in which case the above constraint is met automatically. Otherwise, the efficient (and bilaterally rational) norm can be obtained by letting the above inequality bind, and solving for x_0 . ■