AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Viable Tax Constitutions

Carlo Perroni

And

Kimberley Scharf

No 683

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

THE UNIVERSITY OF
WARWICK

Viable Tax Constitutions*

Carlo Perroni
Warwick University

Kimberley A. Scharf
Warwick University, IFS, and CEPR

Preliminary and Incomplete: First draft January 2003; this draft September 2003

Abstract

Taxation is only sustainable if the general public complies with it. This observation is uncontroversial with tax practitioners but has been ignored by the public finance tradition, which has interpreted tax constitutions as binding contracts by which the power to tax is irretrievably conferred by individuals to government, which can then levy any tax it chooses. However, in the absence of an outside party enforcing contracts between members of a group, no arrangement within groups can be considered to be a binding contract, and therefore the power to tax must be sanctioned by individuals on an ongoing basis. In this paper we offer, for the first time, a theoretical analysis of this fundamental compliance problem associated with taxation, obtaining predictions that in some cases point to a re-interpretation of the theoretical constructions of the public finance tradition while in others call them into question.

KEY WORDS: Taxation, Public Goods, Government.

JEL CLASSIFICATION: H1, H2, H3, H4.

Correspondence should be addressed to Kimberley Scharf, Department of Economics, University of Warwick, Coventry CV4 7AL, UK. E-mail: K.Scharf@warwick.ac.uk

*We are grateful to Tim Besley, Ian Crawford, Ben Lockwood, Hamish Low, David Ulph, and other participants to the 2003 Warwick/IFS Public Economics Weekend for comments and suggestions. All errors are our own.

1 Introduction

The economics literature has traditionally characterized taxation as a coercive levy. The power to tax, in turn—in line with a contractarian view of government institutions which dates back to Hobbes and Locke—can be thought of as arising from a tax constitution to which individuals voluntarily and irreversibly submit in order solve the free-rider problem in public goods provision. A tax constitution would then be the result of a once-and-for-all agreement to confer coercive powers to a monopoly government, a ruling body which pursues goals that are private but nevertheless consistent with the pursuit of collective goals—e.g. surplus extraction, which implies a stake in total surplus—and to which enforcement power is irretrievably surrendered in some way—e.g. by providing the ruler with a strong and unflinchingly loyal army.¹

In modern democracies, however, power is not irretrievably surrendered by citizens to government—arguably, even dictatorships depend on the ongoing support of a (possibly restricted) group of individuals for their survival. If the power of enforcement does not ultimately reside with government but remains within the group that supports government, the notion that the free-rider problem can be solved by a tax constitution seems simply to shift the problem to a different level: how can coercion be sustained in groups?

Given that governments do not possess any independent power of enforcement beyond that of individuals, and given that there exists no independent power of enforcement *outside* the tax constitution, the latter cannot be interpreted as a binding contract in the usual way that we think of contracts between individuals. Rather, enforcement must come from *within* the tax constitution itself. Then, abiding to the tax constitution must be a continuously-renewed, individually and collectively rational choice for individuals and for the group, balancing individuals' temptations for opportunistic behaviour against the group's ability to punish defectors. In other words, a tax constitution must be self-enforcing. This idea is consistent with the point, often made informally, that formal norms are irrelevant if they cannot be enforced, and that policy making must take enforcement problems into account in order to produce relevant policies.

In this paper we develop a theory of self-enforcing tax constitutions under repeated interaction, whereby constrained-efficient tax rules must be supported by renegotiation-proof punishment strategies that are at the same time individually and collectively rational. We then examine how far the idea of a self-enforcing contract can go in helping to rationalize the observed structure of tax institutions, and discuss how it can be reconciled with the presence of political processes presiding collective choices.

The idea that cooperation must be self-supporting has gained currency in the economics literature (e.g. collusion between firms and international agreements) but has never been systematically applied to the interpretation of tax institutions.² Our analysis shows that accounting for the ongoing nature of the incentives for opportunistic behaviour does much more than simply refine old contractarian ideas by giving it a theoretically more satisfying treatment. Sustainability involves stricter requirements on tax structures than contractarianism does—the basic contractarian requirements being that individuals must be better

off under a tax constitution than under a “state of anarchy.” As a result, the analysis of self-enforcing tax constitutions offers sharper predictions about the structure of taxes than contractarianism does. For example, a theory of self-enforcing tax institutions predicts that individuals with the same preferences and income will be pay the same tax (horizontal equity), and that higher income individuals will pay higher taxes (ability-to-pay principle), even if individuals are neutral towards redistribution.

A model of self-enforcing taxation can thus explain the most commonly observed features of tax systems, as traditional models of coercive taxation can, but can derive these features from basic principles of individual and collective rationality—upon which sustainability under repeated interaction is based—rather than from a set of distributional or ethical principles as traditional theories do; and in some cases they go further in offering explanations for certain features that traditional theories are silent on or have difficulty explaining. For example, a model of self-enforcing taxation predicts that there exists a natural upper bound to the level of taxation that can be sustained—something that policymakers seem acutely aware of—an upper bound which lies strictly below income. The axiomatic parsimoniousness of such a model, combined with its ability to account for the problem of enforcement and produce a richer set of predictions, makes it, in our eyes, fundamentally more compelling than traditional theoretical models of coercive taxation.

At the same time, the idea that tax systems reflect a self-enforcing contract calls into question the interpretation of politics as a system of procedures that substitute for bargaining in large groups—an interpretation that is common to much of the recent political economy literature. Instead, a self-enforcing interpretation of tax constitutions points back to welfare theory as a predictive rather than normative tool of tax analysis, and suggests interpreting politics as an information-pooling device rather than as a mechanism for reconciling conflicting interests.

The paper is structured as follows: Section 2 examines constrained-efficient, equilibrium tax constitutions under repeated noncooperative interaction amongst a group of infinitely-lived individuals, supported by renegotiation-proof punishment strategies in which individuals can only be punished through a reduction in collective consumption; Section 3 examines the relationship between collective consumption and property rights, characterizing constrained-efficient, renegotiation-proof tax constitutions in situations where individuals can be also punished by income expropriation but where property rights must also be self-enforcing; Section 4 examines incentive effects when income generation results from an endogenous market effort choice; Section 5 discusses the role of beliefs in punishment; Section 6 focuses on finite lives; Section 7 discusses how the idea of a self-enforcing tax constitution can be reconciled with politics; finally, Section 8 contrasts the idea of self-enforcing cooperation as applied to taxation and collective consumption choices with alternative evolutionary based interpretations of cooperation within groups.

2 Renegotiation-Proof Tax Constitutions with No Income Expropriation

Taxation has traditionally been interpreted as a means of overcoming, through coercion, the free riding problem in public good provision. This problem, cast in modern game-theoretic terms, can be described as follows.

Consider an economy with n individuals, engaging in private and collective consumption. Individual i ($i = 1 \dots n$) receives an income y^i .³ Individuals can contribute to collective consumption an amount c^i , resulting in a level of collective consumption equal to $g = \sum_i c^i$, and in a level of private consumption $x^i = y^i - c^i$. The payoff of an individual in each period is $u^i(x^i, g) = x^i + \theta^i v(g)$, where $v'(g) > 0$ and $v''(g) < 0$ (quasilinear preferences) and $\theta^i > 0$. In a single round of interaction (as examined by, among others, Bergstrom, Blume, and Varian, 1980), a noncooperative equilibrium in contributions will result in a level of collective consumption, $g^N = \sum_i c^{iN}$, that falls short of the efficient level of collective consumption, g^* .⁴ Furthermore, if n is sufficiently large, noncooperative voluntary contributions will be zero.

Coercive taxes can thus be interpreted as an institutional device for supporting collective consumption in large economies where individuals are selfishly-motivated and behave in a rational manner. But, as we have argued in the introduction, unless the power to tax can be irrevocably transferred to an absolute ruler, tax compliance must itself be interpreted as being the result of a deliberate, continuously renewed choice on the part of individuals. Such voluntary compliance can be characterized as an equilibrium phenomenon in a game where strategic interaction is repeated, resulting from a balance between individuals' temptation not to comply and the punishment that the rest of the group could administer against noncompliers.

Once we embrace the idea that tax compliance is voluntary, however, there seems to be little left to distinguish between taxes and voluntary contributions. And indeed the idea that the free riding problem in voluntary contributions to collective consumption can be solved by repeated interaction has been suggested before in the literature (McMillan, 1979). These contributions, like much of the applied literature on repeated interaction, focused on "Nash-reversion" punishment strategies, whereby a reduction in contributions below a certain trigger level in a certain period results in indefinite reversion to the one-shot noncooperative outcome, and concluded that if individuals are infinitely-lived and the above game is repeated indefinitely, levels of contributions above the one-shot noncooperative levels can be supported by the threat of credible punishment in a subgame-perfect noncooperative equilibrium.

There are two fundamental problems with this construction, one conceptual and the other empirical in nature. Nash-reversion punishment strategies are analytically attractive, but there otherwise seems to be no compelling reason to view them as natural candidates for an equilibrium under repeated interaction, and indeed there are good reasons for not doing so. Abreu (1988) has shown that indefinite Nash reversion can be outperformed by

punishment strategies that concentrate punishment in the periods that immediately follow a defection attempt. Furthermore, several writers in the game theory literature have noted that Nash reversion implies threats that are individually but not collectively credible: if a high-contribution outcome can be sustained by the threat of Nash reversion, then, upon entering the punishment phase, all parties would have an incentive to forgo punishment and re-coordinate to a high-contribution equilibrium in the continuation game; this makes the threat of Nash reversion “implausible”. The second objection is related to the first one. Whenever we observe institutions that support cooperation, we never observe indefinite Nash reversion being applied, and this seems also to be the case for voluntary contributions to collective consumption.

To deal with the first type of objection, a number of writers (Farrell and Maskin (1989); Abreu, Pearce, and Stacchetti (1993); Van Damme (1989)) have proposed a further equilibrium refinement beyond subgame perfection for repeated games: namely *renegotiation-proofness*. This imposes a “collective rationality” condition on equilibrium punishment strategies, requiring that, upon entering the punishment phase, it must not be possible to find alternative equilibrium strategies in the continuation game that are favoured by players to the stated course of action; i.e. requiring that one of the parties involved must have an active interest in carrying out the punishment. As Van Damme (1989) has shown with respect to the repeated prisoner dilemma, renegotiation-proofness may require concentrating punishment immediately after defections—as in Abreu (1988), albeit for different reasons.

It turns out that applying the idea of renegotiation proofness to the case of the collective consumption problem also takes care of the empirical objection, because—as we will show—it produces equilibria that more closely resemble the way we observe punishment being actually administered. At the same time, the structure of the resulting equilibria—as we will also show—makes it natural to interpret them as self-enforcing *tax* equilibria.

For the purposes of our discussion, we shall adopt the following characterization of renegotiation proofness: an equilibrium with a level of collective consumption $g < g^*$ is renegotiation proof if it is supported by a continuation equilibrium that is not Pareto dominated by another renegotiation-proof equilibrium featuring the same or a higher level of collective consumption (but a level still short of g^*).⁵ Renegotiation proofness thus requires that the level of collective provision following defection must be rational for nondefectors, so that players could not renegotiate to an alternative, Pareto-dominant equilibrium in the continuation game.

Renegotiation proofness intrinsically embeds the idea of efficient selection not just of the continuation equilibria but also of the candidate equilibrium, since all subgames of the candidate equilibrium, including those on the equilibrium path of play, must be Pareto undominated.⁶ This, however, does not necessarily mean that the level of contributions associated with such an equilibrium must coincide with the maximum sustainable level that can be sustained under renegotiation proofness: with heterogeneous individuals (with respect to income or preferences) and/or differentiated contributions, there may exist equilibria that are undominated by other equilibria featuring higher aggregate contributions, even when collective consumption falls short of g^* . In order to identify equilibria that

support maximum sustainable levels of collective consumption, we need to invoke a further selection criterion. If we consider equilibrium selection according to ex-ante expected utility maximization from the point of view of a representative individual who has yet to observe her own characteristics, quasilinear preferences imply that ex-ante utility uniquely depends on the level of collective consumption, independently of the ex-post distribution of contributions. Thus, equilibrium selection on the basis of the maximum sustainable level of collective consumption—as long as this falls short of g^* —amounts to selection of an ex-ante undominated equilibrium.

Then, under certain conditions (which are discussed below) the following punishment strategies can be shown to be constrained-efficient (i.e. ex-ante undominated), renegotiation-proof equilibrium strategies: each player, i , contributes an amount c^i towards the level of provision $g = \sum_i c^i \leq g^*$, and keeps contributing c^i as long as the other players keep making their respective contributions; if any player i reduces her contribution below c^i in period t , then in period $t + 1$ each of the other players $j \neq i$ reduces her contribution to $\tilde{c}_i^j \leq c^j$, where $\sum_{j \neq i} \tilde{c}_i^j = g_i - y^i$, and where g_i is a collective consumption level, to be defined below, not exceeding g ; this results in a collective reduction in contributions by the nondefectors equal to the difference between g and g_i plus the difference between the defector's income, y^i , and the contribution level that was expected from the defector, c^i ; the other players do so indefinitely until individual i (the defector) increases her contribution to y^i —thus securing a level of collective consumption g_i , after which the other players increase their contributions back to the initial levels. Renegotiation-proofness requires that the level of collective consumption, g_i , experienced by all individuals during the punishment phase be constrained-efficient; such requirement is satisfied if the level of collective consumption at $t + 1$, g_i , and the corresponding contributions, bring about a sustainable, constrained Pareto-efficient allocation, given that the defector is contributing y^i .⁷ The level of provision g_i must be supported at $t + 1$, by analogous punishment strategies against any individual, other than i , who, at $t + 1$, reduces her contribution below the contribution $\tilde{c}_i^j \leq c^j$ required from her during the punishment phase.

The set of constrained Pareto efficient allocations where the defector contributes y^i is defined by contributions $\tilde{c}_i^j \leq y^j, j \neq i$, such that there exist no alternative contributions, $\tilde{c}_i^j \leq y^j, j \neq i$, for which $y^j - \tilde{c}_i^j + \theta^j v(\sum_j \tilde{c}_i^j + y^i) \geq y^j - \tilde{c}_i^j + \theta^j v(\sum_j \tilde{c}_i^j + y^i), j \neq i$, with the inequality being strict for at least one $j \neq i$.⁸

There may be more than one collective consumption level consistent with group rationality as we have defined it, but the punishment strategy that will entail maximum punishment for the defector—and thus induce maximum contributions—will involve adopting a level g_i equal to the lowest level of collective consumption that is consistent with the above collective rationality conditions—a level that will be denoted with g_i^* —if this is sustainable, or the highest level of collective consumption that can be sustained, whichever is smaller.⁹

The conditions for such strategy to be renegotiation-proof equilibrium strategies are then:

$$(c^i - c^{iD}) - \theta^i (v(g) - v(g - c^i + c^{iD})) \leq \delta (y^i - c^i + \theta^i (v(g) - v(g_i))), \quad \forall i, \quad (1)$$

$$(y^i - c^{iR}) - \theta^i (v(g_i) - v(g_i - y^i + c^{iR})) \leq \delta (y^i - c^i + \theta^i (v(g) - v(g_i))), \quad \forall i, \quad (2)$$

$$\sum_{j \neq i} \theta^j (v(g_i) - v(g)) + (g - g_i + y^i - c^i) \geq 0, \quad \forall i, \quad (3)$$

where δ is the one-period time discount factor—the same for all individuals—and where c^{iD} and c^{iR} denote the one-shot unilaterally optimum levels of contribution for player i respectively when the other players contribute levels c^j , $j \neq i$, and when the other players contribute the lower levels $c^j - (g - g_i + y^i - c^i)/(n - 1)$, $j \neq i$.¹⁰

Condition (1) states that the maximum gain for player i in period t from adopting an alternative strategy must be less than the present value of the loss from having to increase her contribution in period $t + 1$.

Condition (2) states that in period $t + 1$, following defection in period t , player i wishes to secure reversion to cooperation from $t + 2$ onwards by paying a higher contribution in period $t + 1$ rather than postponing “repentance” and contributing her unilateral optimum level. It can be readily shown that the maximum sustainable contribution by individual i is indeed attained by making her required contribution in the punishment phase (the fine) as high as possible, i.e. equal to her income, y^i , as implied by (2).¹¹ The fact that the maximum feasible punishment depends on income implies that the maximum sustainable level of collective consumption will vary with income even when it does not exhaust income.¹² It can also be shown that $g_i = \min\{g_i^*, g\}$.¹³ However, if n is sufficiently large and unless preferences are extremely heterogeneous, we will have $g_i = g$.¹⁴ What this means is that, in large economies, the level of collective consumption must remain the same during punishment, and the only plausible threat against defectors is a temporary increase in the contributions required of them. Furthermore, for n large, it will also be the case that both the defection and the no-repentance contribution levels, c^{iD} and $c^{iR} = 0$, will both be zero. For $g_i = g$, the left-hand side of (2) is greater than or equal to the left-hand side of (1), and therefore the latter inequality is always satisfied if the former is.

The left-hand side of condition (3) represents the gain to players other than player i from carrying out the punishment, and is trivially satisfied.¹⁵ This latter condition is what makes such an equilibrium renegotiation-proof.

Result 1 *When income cannot be forcibly expropriated, constrained-efficient renegotiation-proof equilibrium strategies punish a reduction in contributions below their equilibrium-play level by requiring that such a reduction be followed by a temporary increase above the equilibrium-play level.*

The structure of the above renegotiation-proof punishment strategies lends itself to a very natural interpretation: the increase in the contribution required from defectors in the period following defection is consistent with the idea of a *fine*. Fines cannot be applied coercively here (the case where stronger forms of collective coercion can be applied will be discussed later), but are only credible as a way of punishing deviators because of the reduction in collective consumption that results from postponing payment of the fine, and

must be temporary, because the promise of a reversion to the normal state of cooperation is the only way of inducing a deviator to comply with the punishment. Unlike in the case of Nash-reversion punishment strategies, punishing defectors in this way is collectively rational, and hence plausible, as it does not entail a break-up of cooperation.

In such equilibria a decrease in contributions below the “trigger” contribution levels c^i are punished by fines. This feature makes it natural to interpret equilibrium contributions as *taxes*. As we shall see in Section 3, considering stronger means of coercion can give rise to equilibria that can be even more convincingly interpreted as self-enforcing tax equilibria.¹⁶ Nevertheless, suppose that, for the time being, we embrace the interpretation of the equilibrium we have just described as a tax equilibrium. Does this interpretation produce predictions that are in line with the observed structure of taxes?

The first implication we can immediately derive concerns the maximum sustainable level of taxation: a contribution level $c^i = y^i$ cannot be supported, since (2) could not be satisfied in this case; in other words, as long as $g^* > y$, the ratio of contributions to income (i.e., the average tax rate) is bound to be strictly below unity—even if the structure of preferences that has been assumed here does not imply that the ratio of collective to private consumption should be bound below unity—and any increase in disposable income raises the level of collective consumption that can be supported.

Result 2 *The maximum contribution level that can be sustained in a constrained-efficient, renegotiation-proof equilibrium is increasing in the level of income, even when the efficient level of collective consumption is strictly less than aggregate income.*

The application of Nash-reversion punishment strategies produces very different (and implausible) predictions in this respect: with quasilinear preferences, and as long as $g^* \leq ny$ (i.e. the efficient level of collective consumption is an interior optimum given the economy’s resource constraint) the maximum level of contributions that can be sustained by the threat of indefinite reversion to g^N is independent of y ; this implies that supporting a contribution level y may be possible.

If we next look at the relationship between group size and contributions, we find that, for a given per capita income level y , the overall level of contributions that can be sustained in a constrained efficient equilibrium is increasing in the number of individuals involved.¹⁷ On the other hand, holding aggregate income constant, the sustainable level of contributions that can be sustained is decreasing in n , but asymptotically approaches a level that exceeds the one-shot noncooperative level g^N .¹⁸ Unlike with Nash reversion, where the supportable level of collective consumption is independent of numbers, renegotiation-proof punishment is limited by the requirement of collective rationality, and the strength of the punishment that can be plausibly administered decreases with the number of individuals contributing to collective consumption. Thus free riding cannot be fully eliminated. Nevertheless, for a constant level of aggregate income, as n approaches infinity, g approaches a level that will lie, under certain conditions, strictly above the one-shot noncooperative level of collective consumption, g^N .¹⁹ Thus, even in the presence of a renegotiation proofness requirement,

repeated interaction makes it possible to improve on the single-shot outcome, even when the economy is large.

Result 3 *The maximum contribution level that can be sustained in a constrained-efficient, renegotiation-proof equilibrium, for n large, (weakly) exceeds the corresponding equilibrium level of contributions under a single round of interaction.*

If contributions are interpreted in this way, what does the model say about the structure of taxes? As noted earlier, with quasilinear preferences, the ex-ante efficient level of collective consumption is unrelated to the distribution of contributions. Moreover, quasilinear preferences imply no inequality aversion in the model: if we consider equilibrium selection according to ex-ante utility maximization—or, equivalently, take ex-ante expected utility as a measure of social welfare—quasilinear preferences translate into risk-neutrality and hence no inequality aversion. Thus, in this case enforcement constraints are the only determinant of the distribution of contributions in equilibrium.

Nevertheless, when we look at the contributions required from individuals with the same income in a constrained-efficient equilibrium with $g < g^*$, we obtain the predictions that individuals with the same income will pay the same contribution: since the difference between the right- and left-hand sides of (2) is decreasing in c^i , and since these conditions are otherwise the same for all players, the structure of contributions that can best ensure that these inequalities be satisfied for all players involves equal contributions for all players.

Result 4 *When there is no preference heterogeneity in the population, a constrained-efficient, renegotiation-proof equilibrium involves equal contributions by individuals of the same income level.*

When we interpret equilibrium contributions as taxes, the requirement of equal contributions across individuals in a constrained-efficient equilibrium is consistent with the principle of *horizontal equity* in the taxes paid by identical individuals. Here, however, this is just a consequence of the need for supporting a constrained-efficient level of collective consumption, and is fully unrelated to distributional considerations: there is no independent reason for requiring contributions to be the same across individuals other than the need for supporting the highest possible level of collective consumption.

We next look at what the model predicts with respect to how contributions should vary with income. In the discussion that follows, let us first restrict our attention to scenarios where the condition $g_i^* \geq g, \forall i$ is satisfied,²⁰ and focus on the case of a large economy where the aggregate size of contributions relative to individual contributions is such that $\theta^i v'(g - y^i) < 1, \forall i$ so that c^{iD} and c^{iR} are both zero, and where individuals have identical preferences ($\theta^i = \theta, \forall i$) but different incomes. Here, the constrained-efficient contribution levels are increasing in the level of income y^i .²¹ This result is simply a consequence of the fact that the maximum punishment that can be administered against a defector (the fine) depends directly on her level of income.

Result 5 *When there is no preference heterogeneity in the population, an individual's contribution level in a constrained-efficient, renegotiation-proof equilibrium is increasing with her level of income, even when the value of collective consumption to an individual is independent of income and contributions do not exhaust income.*

In the interpretation of contribution levels as taxes, this feature is consistent with the *ability-to-pay* principle in the taxation of individuals with different income levels. Given that preferences are quasilinear in the model, this feature of constrained efficient equilibria is unrelated to distributional concerns. Moreover, the positive relationship between contributions and income levels is not just a consequence of the fact that low income individuals may not have enough income to channel into contributions: constrained-efficient contributions levels will be higher for high-income individuals even when the contribution level required by individual whose income is highest is less than the income of the individual whose income is lowest. Contrasting constrained-efficient, renegotiation-proof equilibria with constrained-efficient, Nash-reversion equilibria, we can note that no positive relationship between contribution levels and income levels is predicted by the latter, absent income effects.

It can also be shown that not only are constrained-efficient contribution levels increasing in income, but they are doing so at an increasing rate.²² The degree of progressivity in a constrained-efficient equilibrium is again unrelated to distributional concerns. Instead, it is directly related to the structure of preferences towards collective consumption—which in the scenario we are currently examining are taken as being common across individuals—i.e. to the shape of the valuation function, v , and to the degree of impatience, which is inversely related to the discount factor δ . Then, other things being equal, contributions will be comparatively less progressive in scenarios where the valuation of the public good provided with the revenue is inelastic (for example, because of the availability of private substitutes to collective consumption) and where individuals are comparatively more farsighted.

Result 6 *When there is no preference heterogeneity in the population, an individual's contribution level in a constrained-efficient, renegotiation-proof equilibrium is increasing more than proportionally to her income level. Equilibrium contributions will be comparatively more progressive the more elastic is the marginal valuation of collective consumption with respect to changes in the level of its provision and the higher the degree of impatience.*

Not only will income inequality result in differentiated contributions; the aggregate level of collective consumption that can be sustained will also depend upon the distribution of income. In particular, progressivity in the maximum sustainable levels of contributions—i.e. convexity of c^i as a function of y^i —implies that an increase in income inequality makes it possible to support more contributions in aggregate.²³ This feature is unrelated to the presence of redistributive goals nor to the aim to offset, through increased collective consumption, the distributional tensions associated with income inequality (as it would be the case, if income can be expropriated—as we show in the next section).

Result 7 *The maximum contribution level that can be sustained in a constrained-efficient, renegotiation-proof equilibrium is larger the more unequal is the distribution of income.*

We can next focus our attention on an alternative scenario where individuals have the same income level but different preferences (different θ^i s). Assume first that preferences are publicly observable. Focusing on a constrained-efficient equilibrium with $g < g^*$, we find that contributions in a constrained-efficient equilibrium are positively related to individual valuations for collective consumption:²⁴ this is simply because, for a given fine, a defector's temptation not to pay the fine is lower the higher is her valuation for collective consumption. This is consistent with the application of the *benefit principle* in the allocation of the cost of provision amongst individuals. It can be also shown, however, that the equilibrium contribution level is less elastic with respect to changes in an individual's valuation for collective consumption than it is with respect to changes in her income level.²⁵ Therefore, even when preferences are fully observable, income differentials have comparatively more effect than valuation differentials in producing a spread in contribution levels, i.e. ability to pay dominates benefit considerations in determining an individual's contribution level in a constrained-efficient equilibrium.

Result 8 *When there is preference heterogeneity in the population, and preferences are publicly observable, an individual's contribution level in a constrained-efficient, renegotiation-proof equilibrium is increasing with her valuation of collective consumption; however, individual equilibrium contribution levels are comparatively more sensitive to changes in income than they are to changes in individuals' valuation for collective consumption.*

Suppose, however, that preferences are not public information and cannot be learned by repeated interaction, so that individual trigger contribution levels cannot be conditioned directly on preferences. This would be the case, for example, if individual preference characteristics, as summarized by the preference parameters, θ^i s, are nonverifiable; then, even if a player's past actions could conceivably be used by others to infer her characteristics, punishment strategies must remain anonymous with respect to the parameter θ^i . An alternative but equivalent assumption would be that the parameters θ^i s are time dependent and serially uncorrelated.

Anonymity does not rule out the use of the punishment strategies we have described, as long as the overall distribution of preference types is fully known. The reason for this is that punishment involves a reduction in contributions by nondefectors, and therefore need not be specifically targeted towards a certain individual: rather than describing the punishment as being triggered by a reduction in an individual's contribution, we can describe it as being triggered by a reduction in the number of contributions of each type below the corresponding equilibrium number. Repentance will then follow on the part of any defector in an equilibrium in which, off the path of equilibrium play, nondefectors stick to the punishment strategies that we have described (i.e. they refrain from increasing contributions themselves in order to induce forgiveness, knowing that the defector will), a behaviour

which is in turn rationally anticipated by the defector. Thus an equilibrium featuring differentiated contributions levels by preference types is fully consistent with anonymity with respect to preference characteristics.

Result 9 *A constrained-efficient, renegotiation-proof equilibrium will feature different contribution levels for different preference types even when preference characteristics are not publicly observable.*

To sum up our discussion so far, if we interpret individual contributions as taxes, in a self-enforcing renegotiation-proof equilibrium where punishment of defectors only consists of a temporary withdrawal of contributions by others, we obtain predictions that are in line with the predictions of models of coercive taxation—such as the prediction that taxes increase with income—as well as other predictions that represent a departure from earlier analyses, and may help explain features of real-world tax systems that traditional analysis are silent upon—such as the prediction that there exist a natural upper bound for the sustainable ratio of collective consumption to income.

Even when predictions parallel those of the traditional analyses, they are derived from different principles. For example, in a constrained-efficient self-enforcing equilibrium, tax progression is not related to inequality aversion but and depends on preferences towards collective consumption. Similarly, the positive relationship between level of tax and benefit received is obtained as a direct result of enforcement constraints rather than derived from the application of ethical principles. Thus, in comparison with traditional theories of coercive taxation, the interpretation of taxes as self-enforcing makes it possible to say more with less.

3 Renegotiation-Proof Tax Constitutions with Income Expropriation

In the collective consumption problem discussed in the previous section, the only recourse for the group against defectors is a reduction in contributions. As shown, once we rule out forms of punishment that are not collectively rational (such as Nash reversion), only limited punishment can be plausibly administered in this way. Nevertheless, we have shown that increasing n does not result in a smaller volume of contributions.

Consider, however, the following alternative conceptual experiment. Suppose that we start at a certain constrained-efficient equilibrium where collective consumption is g now suppose that we increase the number of individuals from n to $n' > n$ and simultaneously change the per-capita income and level to $y' = (n/n')y$ and change the public good valuation scale parameter from θ to $\theta' = (n/n')\theta$. In the resulting modified economy, aggregate income is the same as in the original economy, and the aggregate valuation for collective consumption also remains the same. The efficient level of collective consumption is constant. The per-capita level of contribution required to support a given level g , however, is now

smaller, and so is per capita income; then, the expressions $y - c$ and $\theta v(g)$ on the left-hand side of (2) and $y - c$ on the right-hand side are all proportionally scaled down; the expression $\theta v(g - y)$, however, decreases less than proportionally, implying that (2) will be violated in the modified economy. Thus, g will approach zero as we continue scaling up the economy. We should then conclude that the scope for funding collective consumption through self-enforcing taxes is limited in a large economy: although repeated interaction can make it possible to improve on the one-shot noncooperative outcome, the free-riding problem remains.

Furthermore, interpreting equilibrium contributions in that model as representing taxes supported by “voluntary” fines also runs against the observation that stronger forms of coercion are available and are actually used in support of tax compliance. However, for us to be able to rationalize the use of stronger forms of punishment in the context of a self-enforcing equilibrium, it must be the case that their application is ex-post rational and hence ex-ante credible. For example, punishing a reduction in contributions by death is not credible if administering the penalty is costly, simply because applying the penalty is not an ex-post rational response for nondefectors (neither individually nor collectively) once the defector adopts her one-shot optimal level of contribution.

One channel through which punishment can be credibly administered is income expropriation, if it is feasible. Expropriation can be credible (both individually and collectively) because the punishers stand to gain from expropriating a defector’s income. However, expropriation adds a new source of temptation to behave opportunistically; thus, if we call expropriation into play as a form of punishment against defectors, we must also account for an economy’s ability to prevent expropriation attempts from occurring in equilibrium. For this to be possible, there must be force in numbers: a group of individuals acting against a defector must be able to confiscate from the defector an amount that is greater than the total amount an individual can gain by individually attempting expropriation against members of that group. This idea can be formalized as follows. Each individual is endowed with a certain amount of time—normalized to unity without loss of generality—which can either be used either for income generation, obtaining a maximum income y^i , or to engage in expropriation attempts against others up to a total amount of time representing a fraction ρ of total available time. Only a fraction ζ of income, representing income generated from market activities, can be expropriated, with the rest of income representing nonmarket income.²⁶ An expropriation attempt by i against an individual j uses a fraction $\gamma < \rho$ of i ’s time, and therefore involves a cost γy^i . The individual can thus at most participate in $\bar{n} = \rho/\gamma$ expropriation attempts. (For simplicity, we shall assume $\bar{n}$ to be an integer number.) If a group of individuals other than i , consisting of $s \leq n$ individuals, attempt expropriation against j , they can collectively expropriate a fraction, $\beta(s) = \xi(s)\zeta$, of her income, y^j —obtaining an amount equal to $\beta(s)y^j$ —where $\beta(s) = \xi(s)\zeta$, $0 \leq \xi(s) \leq 1$, $\xi'(s) > 0$, $\xi''(s) < 0$.²⁷

In this framework, property rights over transferable goods are not exogenously enforced, and can only emerge as an equilibrium phenomenon (when the equilibrium level of expropriation effort is zero). On the other hand, initial possession of income entails a preferential

claim because forcible expropriation is costly, so that efficient outcomes will feature as little expropriation as possible: aggregate disposable income will be maximized when no expropriation attempts take place. But if $\beta(1) \equiv \alpha$ is large enough and γ is not too large (i.e. if expropriating others' income is sufficiently easy), an efficient, no-expropriation outcome cannot be supported within a single round of interaction. To see this, focus on the symmetric case where all individuals are identical ($y^i = y, \forall i$). In this case, if $\alpha > \gamma$, a symmetric, pure-strategy, no-expropriation outcome does not constitute a one-shot noncooperative equilibrium. Repeated interaction, on the other hand, may make it possible to prevent expropriation attempts from occurring.

What we are primarily concerned with for the purposes of our discussion is the question of how the possibility of income expropriation affects equilibria with collective consumption. Before examining this question, however, we shall first look at expropriation in isolation from the collective consumption problem.

3.1 Equilibrium Property Rights

Suppose that income is only used for private consumption ($\theta = 0$). Under certain conditions (stated below), an efficient outcome where no expropriation attempts take place may be supported by renegotiation-proof punishment strategies in which punishment only lasts for one period, and which are described as follows (it will be notationally convenient in what follows to let individuals be ordered by the size of their potential income, i.e. $y^{i+1} \geq y^i, \forall i$): each player i , $i \in I$, does not attempt expropriation against any other player as long as the other players do the same; if any player i deviates from this course of action at t , then in period $t + 1$ player i becomes an expropriation target for a group of $s^i \leq n$ of the other players and fully accommodates expropriation by others; if player i accommodates punishment in this way at $t + 1$, all players revert to no expropriation from $t + 2$ onwards, otherwise forgiveness of player i is postponed until player i accommodates punishment.²⁸

Focus on a scenario where unilateral expropriation is always attractive to all individuals ($\alpha y^j \geq \gamma y^i, i \neq j$). If player i attempts expropriation, the highest gross deviation gain she can obtain is by directing her expropriation attempts against the $n'_i \leq \bar{n}$ highest income individuals—other than herself—for which $\alpha y^j \geq \gamma y^i, j \geq n - n'_i$. Denote the set of such individuals as S_R^i ; the net gain, $\Lambda^i \geq 0$, that can be obtained from these unilateral expropriation attempts is

$$\Lambda^i \equiv \sum_{j \in S_R^i} \max\{\alpha y^j - \gamma y^i, 0\}. \quad (4)$$

Λ^i will be decreasing in i (i.e. as we move up the income distribution) may reach zero at a certain income level (and remain zero thereafter). In other words, attempting expropriation will be relatively more tempting for low-income individuals than for high-income individuals.

Assuming that n is large, so that $\rho/\gamma < n$, the punishment strategy just described will

be a renegotiation-proof equilibrium strategy if

$$\Lambda^i \leq \delta \beta(\bar{s}^i) y^i, \quad \forall i; \quad (5)$$

where the set of players, $S_P^i(\bar{s}^i)$, participating in expropriation of the defector in the punishment phase consists of the largest possible number, $\bar{s}^i$, of lowest-income individuals, other than player i , who can collectively profit from carrying out the punishment, i.e. the largest s_i for which

$$\beta(s^i) y^i - \gamma \sum_{j \in S_P^i(s^i)} y^j \geq 0. \quad (6)$$

Condition (5) ensures that the gain to player i from deviating and/or from postponing repentance after deviation must be less than the present value of the punishment.³⁰ Condition (6) says that expropriating the defector must produce a net benefit for nondeviators in comparison with the no-expropriation path.³¹

In an economy where the number of individuals of each income type is greater than both ρ/γ and s^i for any possible value of s^i satisfying (6), we have $\Lambda^i = (\rho/\gamma) \max\{\alpha y^n - \gamma y^i\}$, $\forall i$, and $y^i = y^1$, $i \leq s^i$. Let $\bar{\beta}^i$ denote the maximum collectively rational level of expropriation against individual i , i.e. $\bar{\beta}^i = \beta(\bar{s}^i)$, where $\bar{s}^i$ is the highest value for which $\beta(\bar{s}^i) y^i - \gamma \bar{s}^i y^1 \geq 0$. Note that $\beta''(s) < 0$ implies that $\beta(s)/s$ is decreasing in s , and therefore $\bar{s}^i$ is increasing in y^i . This means that the difference between the right- and the left-hand sides of (5) will always be smallest for the individuals with the smallest income, i.e. for $i = 1$.

A no-expropriation outcome—where property rights are established—can then be supported in equilibrium in a large economy if

$$\delta > \frac{\rho}{\bar{\beta}^1} \left(\frac{\alpha}{\gamma} \frac{y^n}{y^1} - 1 \right). \quad (7)$$

Condition (7) implies that for property rights are easier to establish the “weaker” individuals are relative to groups as credible expropriators, i.e. the smaller is $\alpha\rho/\gamma = \alpha\bar{n}$ relative to $\bar{\beta}^1 > \alpha$. For (7) to be satisfied when players are impatient ($\delta < 1$), we must have

$$\frac{\alpha\rho}{\gamma} \frac{y^n}{y^1} - 1 < \bar{\beta}^1. \quad (8)$$

For a given α , the right-hand side of (8) is decreasing in both γ (the cost of expropriation) and in the ratio $\bar{\beta}^1/\alpha$ (the “strength” of the group relative to that of an individual). A larger γ , however, reduces the profitability of expropriation for the punishers, which can result in a reduction in $\bar{\beta}^1$. Therefore, establishment of property rights is more likely if the ratio $\bar{\beta}^1/\alpha$ is comparatively large and γ is comparatively small. Consider, for example, the case of a developed economy, where income originates from market activities and consists of marketable output, and contrast this with the case of a less developed economy, where income stems mostly from nonmarket activities; in the first case $\bar{\beta}^1$ will be comparatively large and γ comparatively small (expropriation of marketable output is comparatively easy) and therefore it will be comparatively easier for property rights to be established. When y^i is the same for everyone, (8) requires $\alpha/\gamma < 1 + \bar{\beta}/\rho$.³²

Result 10 *The establishment of property rights requires expropriation costs to be comparatively small and group strength to be comparatively high.*

Condition (7) also says that property rights are easier to establish the larger is the gap between mean income and lowest income: thus, inequality poses a threat to property rights; and if income inequality is sufficiently large, it will never be possible to sustain a no-expropriation equilibrium for $\delta \leq 1$.

Result 11 *The more unequal is the distribution of income, the more patient players must be in order for a no-expropriation outcome to be sustainable under repeated interaction.*

If direct transfers between players are feasible and are used along the equilibrium path, then it may make it possible to prevent expropriation from occurring even when incomes are unequal. Suppose that it is possible for an individual i to transfer voluntarily to individual j an amount $r^{ij} \geq 0$. Such transfers will never be used as part of a one-shot noncooperative equilibrium, but they may be used under repeated interaction as part of an equilibrium where the withdrawal of transfers is used as a credible punishment against defections. Then condition (5) becomes

$$\Lambda^i + \sum_j r^{ij} \leq \delta \left(\bar{\beta}^i y^i + \sum_j r^{ji} \right), \quad \forall i. \quad (9)$$

A simultaneous increase in the transfer made and received by an individual increases the deviation gain relative to the punishment for that individual; we can thus restrict our attention to equilibria where individuals either receive transfers from others or make transfers to other, i.e. where $\sum_j r^{ij} \sum_j r^{ji} = 0, \forall i$. Let c^i denote the aggregate net transfer, made or received by individual i : $c^i = \sum_j r^{ij} - \sum_j r^{ji}$. Then for n large, the above two inequalities can be written as

$$\frac{\alpha\rho}{\gamma} y^n - \rho y^i + \max\{c^i, 0\} \leq \delta \left(\bar{\beta}^i y^i + \max\{-c^i, 0\} \right). \quad (10)$$

The fact that the net transfer c^i enters these conditions asymmetrically depending on its sign is a consequence of the fact that while a positive c^i (a positive transfer by i) is the result of a choice by i (and therefore belongs to i 's strategy space), a negative c^i (a transfer received by i) is the result of others' choices; hence c^i goes to zero in the context of a deviation by i only if positive, and goes to zero as part of i 's punishment only if negative. It may then be possible to find values $c^i, \forall i$, satisfying $\sum c^i = 0$, such that the above inequalities are satisfied for all individuals. Transfers can ease distributional tensions by increasing the opportunity cost of defections for low-income individuals—the individuals who are most tempted to expropriate.

Result 12 *Transfers (negative taxes) may be needed to support a no-expropriation equilibrium.*

Those transfers can be interpreted as representing a system of redistributive taxes, which are not the result of preferences for redistribution but serve to complement group expropriation as a form of punishment against defectors in support of property rights. Even if the potential victims of expropriation must surrender part of their income to the potential expropriators, these equilibria Pareto dominate equilibria with expropriation because they avert expropriation costs.

On the other hand, redistributive transfers may not be enough for a no-expropriation equilibrium to be sustainable, as the following example shows. Suppose that half of all individuals have a certain income y' and the other half have an income of zero. With n large, the net gain a zero-income individual gets by expropriating others is then $(\alpha\rho/\gamma)y'$. In order to be induced not to expropriate, the zero-income individuals must each receive a net transfer greater than $(\alpha\rho/\gamma)y'/\delta$; then, since the number of high- and low-income individuals is the same, the cost of funding such transfers to each one of the high-income individuals would also be equal to $(\alpha\rho/\gamma)y'/\delta$. Therefore, if $1 > \alpha\rho/\gamma > \delta^2\beta(\bar{s}')$, the defection gain for a high-income individual is larger than the punishment incurred;³³ in this case, a no-expropriation equilibrium cannot be supported even if transfers are feasible.

3.2 Income Expropriation and Collective Consumption

Having dealt with the question of how property rights can be sustained, we are now in a position to ask how the feasibility of income expropriation affects the sustainability of private contributions to collective consumption, in an economy where collective consumption has a positive value ($\theta > 0$).

When expropriation is possible, the following punishment can be used in place of the punishment strategy described in Section 2: if an individual i deviates from equilibrium play in period t , the other individuals expropriate i in period $t + 1$, extracting from i an amount equal to $\beta^i y^i \leq \bar{\beta}^i y^i$, which is directed to fund collective consumption, and in the same period they collectively reduce their contributions by an amount equal to $g - g_i + (\beta^i + \eta^i(1 - \beta^i))y^i$, $0 \leq \eta^i \leq 1$, (the difference between g and their collectively rational level of contributions plus the amount of income expropriated from i plus a fraction of the portion of i 's income that is not expropriated); player i accommodates expropriation at $t + 1$ and at the same time changes her contribution to $\eta^i(1 - \beta^i)y^i$; if player i does this, then all players revert to equilibrium play from $t + 2$ onwards, otherwise forgiveness is postponed until player i accommodates punishment. This punishment strategy (weakly) dominates the no-expropriation strategy described in Section 2 in terms of the punishment that can be inflicted on defectors, since the temptation not to accommodate punishment by making contributions is comparatively reduced: the contribution required by the defector in the punishment phase is now $\eta^i(1 - \beta^i)y^i$ rather than y^i .³⁴

Linked punishment will always be used in a full-information constrained-efficient equilibrium: deviation in contributions and expropriation choices, whether occurring in isolation from each other or jointly, will be best prevented by relying on the strongest available punishment. Given this, the deviations that we need to consider—the deviations that generate

the largest one-shot gains for deviators—are joint deviations.

The expropriation-based punishment strategies we have described are only actually credible (subgame perfect) as part of an equilibrium where no deviations occur, thus where potential deviations are by single individuals only. Given the technological constraints associated with expropriation, a joint deviation by a subset of individuals could not actually be punished in this way.

For simplicity, throughout the following discussion we restrict our attention to scenarios where the number of individuals is sufficiently large and/or individual valuations are sufficiently small relative to the aggregate valuation for collective consumption that one-shot optimal deviations always involve zero contributions, i.e. where c^{iD} and c^{iR} are both zero. Also assume for now that the collectively rational, sustainable level of contributions for all players other than i , g_i is equal to g . Then, the conditions for the above punishment strategy to be a renegotiation-proof equilibrium strategy can be written as

$$\begin{aligned} & \max\{c^i, 0\} - \theta^i (v(g) - v(g - \max\{c^i, 0\})) + \Lambda^i \\ & - \delta ((\beta^i + \eta^i(1 - \beta^i))y^i - c^i + \theta^i (v(g) - v(\min\{g, g_i^*\}))) \leq 0, \end{aligned} \quad (11)$$

$$\begin{aligned} & \eta^i(1 - \beta^i)y^i - \theta^i (v(\min\{g, g_i^*\}) - v(\min\{g, g_i^*\} - \eta^i(1 - \beta^i)y^i)) + \Lambda^i \\ & - \delta ((\beta^i + \eta^i(1 - \beta^i))y^i - c^i + \theta^i (v(g) - v(\min\{g, g_i^*\}))) \leq 0, \end{aligned} \quad (12)$$

$$\sum_{j \neq i} \theta^j (v(\min\{g, g_i^*\}) - v(g)) + g - \min\{g, g_i^*\} + \beta^i y^i - \gamma \sum_{j \leq s^i, j \neq i} y^j \geq 0. \quad (13)$$

Result 13 *When income can be forcibly expropriated, constrained-efficient renegotiation-proof equilibrium strategies punish a reduction in contributions below their equilibrium-play level by a temporary increase above the equilibrium-play level in the defector's payment, part of which is obtained by direct expropriation and part of which is voluntarily contributed by the defector.*

Contributions that are supported by a combination of fines and forcible income expropriation even more closely resemble real-world taxes. So, from now on we shall refer to them as taxes.

If the defector were to voluntarily pay the whole penalty, y^i , the group would manage to save the costs associated with forcible expropriation; so it would be possible, in principle, to find an alternative outcome that Pareto dominates the one where punishment relies on expropriation. However, since the repentance constraint is the binding constraint without expropriation, and since expropriation punishment relaxes that constraint, if the repentance constraint with expropriation binds (at a constrained-efficient optimum) all alternative continuation equilibria that feature less expropriation against the defector will involve lower taxes from the defector in the following periods.³⁵ Expropriation will therefore always be relied upon by constrained-efficient punishment strategies. On the other hand, if for a

certain value $\beta^{i'} < \bar{\beta}^i$ the no-defection condition also becomes binding, then raising β^i above $\beta^{i'}$ cannot be part of a renegotiation-proof equilibrium, since in that case it would be possible to re-coordinate to an alternative equilibrium featuring less expropriation (given that in that case the first constraint is the binding one) and no less collective consumption: if $\eta^i = 1$, the first constraint is independent of β^i , and if $\eta^i < 1$ it is possible to lower β^i and increase η^i so as to relax the first constraint.

It can be shown that the maximum supportable tax for a certain individual i in a constrained-efficient, renegotiation-proof equilibrium is identified by one of three possible regimes: (i) the repentance constraint is the binding constraint, the maximum supportable level of expropriation is applied against i in the punishment phase ($\beta^i = \bar{\beta}^i$), and i voluntarily pays all of her residual income in the punishment phase ($\eta^i = 1$); (ii) both the no-defection and the repentance constraints are binding, the maximum supportable level of expropriation is applied against i in the punishment phase ($\beta^i = \bar{\beta}^i$), and i voluntarily pays a fraction of her residual income in the punishment phase ($\eta^i < 1$); (iii) both the no-defection and the repentance constraints are binding, the level of expropriation applied against i in the punishment phase falls short of its maximum supportable level ($\beta^i < \bar{\beta}^i$), and i voluntarily pays a fraction of her residual income in the punishment phase ($\eta^i < 1$).³⁶ Regime (i), where β^i and η^i are at their respective upper bounds, is analogous to the no-expropriation regime of Section 2, in that the sustainable tax is identified here uniquely by the repentance constraint; the overall level of payment by a defector during the punishment phase will also in this case be y^i , but expropriation will be relied upon as a partial substitute for voluntary fines, as it relaxes the repentance constraint and raises the sustainable tax. In regime (ii) expropriation substitutes for voluntary fines, which are not used to their full available extent. Regime (iii) represents a situation where the ability of the group to credibly punish defectors with expropriation exceeds what can be plausibly promised, given that applying maximum expropriation would violate collective rationality (since it would be unnecessarily wasteful). In this latter regime, expropriation effectively makes the repentance constraint irrelevant; it is as though, in a no-expropriation scenario, the no-defection constraint were the binding one (given that β^i and η^i do not effectively appear in the first constraint).

Which regime a certain individual i falls under will depend on her income level as well on the other parameters of the problem, and particularly on the size of $\bar{\beta}^i$. Consider the following experiment: starting from a scenario where $\bar{\beta}^i = 0$, suppose that expropriation against i during punishment becomes progressively more available, i.e. the upper bound $\bar{\beta}^i$ increases. For $\bar{\beta}^i$ small, the second-constraint will remain the binding constraint and its derivative with respect to η^i will remain negative, implying that η^i will be at its upper bound and $\beta^i = \bar{\beta}^i$ (regime (i)); however, as we increase $\bar{\beta}^i$ further, the derivative of the second constraint with respect to η^i will become positive and we will move to a regime where $\eta^i < 1$ and $\beta^i = \bar{\beta}^i$ (and both constraints are binding) (regime (ii)); if $\bar{\beta}^i$ becomes even larger, we move to regime (iii), where the maximum available level of expropriation exceeds the level that will be relied upon in a constrained-efficient renegotiation-proof equilibrium. In regimes (i) and (ii), a change in $\bar{\beta}^i$ affects the sustainable tax. This implies that, if $\bar{s}^i$ is

increasing with y^i , the possibility of relying on expropriation during punishment will have different effects for individuals with different incomes. In regime (iii), on the other hand, different $\bar{\beta}^i$ s for different individuals will have no effect. Nevertheless, since s^i is increasing in y^i , higher-income individuals will be more likely to fall within regime (iii) than lower income individuals.

Relying on expropriation during punishment strictly dominates relying only on increased contributions as a punishment device. However, expropriation also increases the temptation to defect by expropriating others. With expropriation it may be impossible to support any positive level of tax for low income individuals for whom $\Lambda^i > 0$ in a constrained efficient equilibrium featuring no expropriation. In such situations, a constrained efficient equilibrium may require transfers (negative taxes) to low-income individuals, as discussed in the previous subsection. For these individuals, the negative tax, unlike a positive tax, will remain unchanged during deviations and will vanish in the punishment phase. This means that c^i does not appear on the left-hand side of the no-defection condition, and the repentance constraint will always be (weakly) binding.

Since expropriation also increases the temptation to defect, both in the cooperation phase and in the punishment phase, the effect of expropriation on an economy's ability to sustain a cooperative outcome is generally ambiguous. However, if we focus on a scenario with identical individuals, we can show that in economies where property rights can be independently established—i.e. in economies where condition (7) is satisfied— expropriation possibilities always help to solve the free-riding problem in the provision of collective consumption.³⁷ Thus, if expropriation goes hand in hand with sustainable property rights, it can make it possible to support higher levels of collective consumption.

Result 14 *The maximum level of collective consumption that can be supported in an economy where expropriation is feasible and where property rights are otherwise established is higher than the corresponding level in an economy where expropriation is not feasible.*

If we repeat the scaling experiment we described at the beginning of this section—whereby y and θ are both scaled down as n increases—we can now see that, with expropriation, g will approach the positive asymptote $(1 - \bar{\beta}/\delta)Y$ (as long as $\delta > \bar{\beta}$). Thus, expropriation can make it possible to overcome the increase in free riding associated with large numbers, in the sense that it will make it possible to support a level of collective consumption that does not vanish in size, when measured relatively to the efficient level of collective consumption, as n becomes progressively larger.

Result 15 *In a large, homogenous economy where income expropriation is feasible, holding total income and the optimal level of collective provision constant, an increase in the number of individuals leaves the maximum sustainable level of collective consumption in a constrained-efficient equilibrium unchanged.*

Our analysis also points to the existence of a relationship between property rights and collective consumption. The enforcement of property rights has been recognized as one of

the essential activities of government, and has even been described by early public finance writers as a public good. Technically, the enforcement of laws and contracts cannot be classified as a form of nonrivalrous collective consumption, as single acts of enforcement benefit specific individuals; rather the collective dimension of enforcement consists of its reliance on the strength of the group vis-à-vis the individual. Nevertheless public finance writers and economic historians have long recognized that law enforcement and the provision of collective goods have always co-existed since the birth of government. The theoretical framework we have just described enables us to explore this relationship formally.

As noted earlier, we can think of expropriable output as being associated with market activities. For given characteristics of the expropriation technology, as summarized by the function $\xi(s)$, what is the effect on the supportable level of collective consumption of an increase in the relative importance of market activities if this simply translates in an increase in the portion of income that can be expropriated? In our model, this amounts to examining the effects of an increase in ζ , resulting in an equiproportional upward shift in the expropriation schedule $\beta(s) = \xi(s)\zeta$.

It can be shown that, in regimes (i) and (ii), where group expropriation possibilities are fully exploited, (7) is a sufficient condition for an increase in ζ to result in an increase in g . We can interpret this result as predicting that the emergence of market activities can facilitate the development of collective consumption activities.

Result 16 *In a large, homogenous economy where income expropriation is feasible, an increase in the share of income that can be expropriated can raise the maximum sustainable level of collective consumption.*

This is consistent with the observation that collective consumption is a feature of advanced market economies, a feature that has been typically associated with the idea that the collective provision of infrastructure may be required for the development of market activities. While this form of causation is undoubtedly present, our analysis shows that the reverse form of causation may also be important.³⁸

When we consider regime (iii), however, a positive relationship between ζ and g is no longer guaranteed. Even when (7) is met, an increase in ζ which causes a switch to regime (iii) can result in a decrease in g ; this is because in regime (iii) group expropriation is constrained not by technology but by the requirements of renegotiation proofness, so that the only effect of an increase in the share of expropriable income is a stronger temptation to expropriate others, which in turn reduces the maximum tax that can be sustained. Thus as the share of income that can be expropriated progressively increases, the maximum sustainable level of collective consumption can first increase and subsequently decrease. The latter pattern would be representative of a “mature” economy, where most income is derived from market activities.

Conversely, it can be shown that the need for engaging in collective consumption can facilitate the establishment of property rights in situations where it is not possible to do so otherwise. Specifically, in situations where transfers are required to prevent unilateral expropriation attempts from occurring, channeling the transfers through collective

consumption can make it easier to enforce property rights. This can be shown by considering the following example. Let us go back to the scenario introduced earlier, where one-half of the population has income y' and the other half have an income of zero, and suppose that preferences are common across consumers ($\theta^i = \theta, \forall i$). Suppose also that $1 > \alpha\rho/\gamma > \delta^2\beta(s-1)$, and therefore (as discussed earlier), a no-expropriation equilibrium cannot be supported by the use of transfers in the absence of collective consumption. The level of collective consumption provision which maximizes aggregate (or expected) utility in this case is $g^*(n)$ s.t. $n\theta v'(g^*(n)) = 1$, but the minimum value of collective consumption that is consistent with Pareto efficiency (given the constraint that zero-income individuals cannot contribute to collective consumption or otherwise compensate contributors) is $g^*(n/2) < g^*(n)$ s.t. $(n/2)\theta v'(g^*(n/2)) = 1$. Consider an equilibrium with $g = g^*(n)$; then, the punishment to a zero-income individual for attempting expropriation consists of a temporary fall from $g^*(n)$ to $g^*(n/2)$ in the level collective consumption experienced. Then, in order for a zero-income individual to be persuaded not to expropriate, we must have that $(\alpha/\gamma)y' \leq \delta\theta(v(g^*(n)) - v(g^*(n/2)))$, and for a high-income individual to be persuaded to contribute $2g^*(n)/n$, we must have $2g^*(n)/n - \theta(v(g^*(n)) - v(g^*(n) - 2g^*(n)/n)) \leq \delta(y' - 2g^*(n)/n + \theta(v(g^*(n)) - v(g^*(n/2))))$. Suppose the first condition is satisfied with equality. Then, for n large, the second condition can be rewritten as $\alpha/\gamma > (1 + \delta - \theta v'(g^*(n)))(g^*(n)/(ny'/2)) - \delta$. One can verify that this condition is compatible with the condition $1 > \alpha/\gamma > \delta^2\beta(s-1)$, implying that taxes that are used to fund collective consumption can manage to prevent expropriation from occurring in situations where direct transfers cannot.

The reason for this result is that, although the threat of reduced provision of collective consumption can deter expropriation attempts just as the threat of reduced transfers can, indirect transfers are easier to sustain for those making them if the latter also benefit from the increased collective consumption, both because this lowers the temptation to temptation to withhold the transfer, and because it increases the punishment that other contributors can administer to those who fail to contribute. Furthermore, the value of an indirect transfer can be larger to the receivers than that of a direct transfer if collective consumption is at a suboptimal level for which the sum of the receivers' marginal valuations is greater than the cost; i.e. the fall from $g^*(n)$ to $g^*(n/2)$ could be worth more to the low-income individuals than the actual money transfer associated with it, for if they received a direct transfer instead, they would themselves be unable to sustain a level of collective consumption $g^*(n)$, because the punishment that can be administered against them is limited. In this case, the high-income individuals vicariously solve the free-riding problem for the zero-income individuals through indirect transfers.

As an extreme illustration of this mechanism, suppose that only the low-income people value public goods, but are unable to sustain collective provision themselves because of their low income. Higher-income individuals are then indifferent between making direct transfers and contributing to collective consumption; contributions towards the public good, however, are worth more to the low-income recipients, and the plausible threat of terminating them is therefore comparatively more effective as a deterrent against expropriation attempts by

low-income individuals. This feature is also consistent with governments' use of public good provision as a redistributive tool—to which a sizable theoretical literature has been devoted.

Result 17 *The need to engage in collective consumption can make it possible to sustain property rights even in situations where it is not otherwise possible to do so.*

Suppose, however, that only direct transfers can make it possible to support property rights. Then a constrained-efficient equilibrium may feature transfers even in the presence of taxes funding collective consumption ($\theta > 0$). This would trivially be the case, for example, in a scenario where half of the population have no income ($y'' = 0$) and where zero-income individuals do not value collective consumption ($\theta' = \theta > 0, \theta'' = 0$). Thus, redistributive provision of collective goods can co-exists with direct redistributive transfers.

Given that collective consumption, used as an indirect transfer, can be more effective at preventing expropriation, we could expect that more income inequality, which increases the temptation of unilateral expropriation attempts, may lead to increased collective consumption. However, the maximum tax that is sustainable for high income individuals is independent of the expropriation gain by low income individuals; thus, while collective consumption can be effective in preventing expropriation attempts from occurring, a higher threat of unilateral expropriation by the poor does not translate into higher taxes—which are only a function of the plausible threat of collective (rather than unilateral) expropriation being carried out as punishment. On the contrary, an increase in income inequality may trigger a switch from a situation where low-income individuals are tempted to expropriate higher income individuals to one where the temptation is positive; then, if direct rather than indirect transfers are more effective at preventing expropriation attempts (say, because the potential expropriators do not value collective consumption), increased inequality would require direct transfers, which, for a given sustainable level of tax revenues, would reduce the level of collective consumption that can be funded. Increased inequality, however, can raise collective consumption through a different channel: other things equal, more inequality can increase the gain from collective expropriation of defectors, because it lowers the opportunity cost of expropriation relative to the income expropriated, and can therefore result in a more severe punishment of defectors—the maximum expropriation level, $\bar{\beta}^i$, that is consistent with renegotiation proofness increases. Consider, for example, an economy where $\alpha = 0$, implying that there is no temptation to attempt expropriation individually, but where $\beta(s) > 0$ for $s > 1$, and where one-half of the population has a certain income level and the other half a lower income level. Then, a mean-preserving income spread would leave the maximum level of collective expropriation against poor defectors in the punishment phase unchanged, but would raise the maximum level of collective expropriation against rich defectors, and total tax revenues would then (weakly) rise as a consequence.

Result 18 *In an economy where income expropriation is feasible, a mean-preserving income spread has an ambiguous effect on the maximum sustainable level of collective consumption.*

If we next examine the question of how taxes vary with income, we also find that, as in the no-expropriation scenario discussed in Section 2, individuals with the same income will pay equal taxes (horizontal equity) and contributions will be increasing with income (ability to pay). As noted, however, taxes can be negative in the presence of expropriation. Furthermore, the rate of tax progression in this case will also depend on the expropriation technology, and will be different depending on whether taxes are positive or negative and depending on the optimality regime that individuals fall under.³⁹ Other things equal, marginal tax rates will be higher for low-income individuals who are potential expropriators ($\Lambda^i > 0$) (and, in particular, for individuals who are recipients of transfers) than for higher-income individuals who are not: this is because a marginal increase in income for a potential expropriator directly reduces the temptation to engage in expropriation; this relaxes the relevant enforcement constraint(s), which in turn makes it possible to further raise the tax required from the individual. Such effect is absent in the case of high-income individuals who face no expropriation incentives ($\Lambda^i > 0$). Note, however, that in the case of a tax schedule that features negative taxes for low income individuals, falling marginal tax rates can co-exist with increasing average rates.

Result 19 *In a constrained-efficient, self-enforcing tax equilibrium, marginal tax rates will be positive. Ceteris paribus, negative taxes will be associated with comparatively higher marginal tax rates.*

Real-world tax systems do rely on coercive expropriation, and feature negative taxes at the bottom end of the income distribution characterized by high replacement rates (progression); also, as predicted by our analysis, progressivity typically flattens out at the top of the income distribution—a feature that traditional analyses of optimal taxation have been at pains to explain.⁴⁰

3.3 Voluntary Contributions

If taxes are simply interpreted as self-enforcing contributions—albeit supported by the threat of income expropriation—then how do we explain the existence of “voluntary” contributions, contributions which individuals view as being separate from taxes and noncompulsory? It turns out that unobservable heterogeneity, combined with observable income differentials, may result in equilibria supported by punishment strategies that only partially rely on expropriation, and where taxes and “voluntary” contributions can be separately identified, as the following discussion illustrates.

Suppose that there is preference heterogeneity in the economy and that preferences are not observable, and focus on the case where the income level is the same for everyone. There can then be three possibilities for a constrained-optimal equilibrium: (a) everyone is required to pay a common contribution level (tax) supported by the punishment strategies described in the previous subsection, and there is full compliance; (b) everyone is required to pay a common contribution level (tax) supported by the same punishment strategies,

but there is only partial compliance, with lower-preference individuals not paying the tax; (c) everyone pays a certain minimum contribution level (tax) supported as above, but, in addition, a subset of players make a contribution above the tax, with these additional contributions being supported by the punishment strategies described in Section 2—anonymous punishment by reduction in collective contributions only, without invoking the threat of expropriation. An equilibrium such as that described under (c) will always (weakly) dominate an equilibrium as under (a), since using the anonymous punishment strategies in addition to the expropriation-based strategies can never result in lower aggregate contributions. An outcome such as (b) may be favoured if the number of high-preference individuals relative to that of low-preference individuals is sufficiently large that inducing the highest possible level of contributions by them—by relying on the threat of expropriation—dominates the cost of defection by the low-preference individuals.

For example, suppose that all individuals have the same income but there are $\underline{n}$ individuals with $\underline{\theta}$ and $\bar{n}$ individuals with $\bar{\theta} > \underline{\theta}$. Also suppose that individuals are not tempted to expropriate others ($\alpha = 0$) and group expropriation possibilities are not binding (regime (iii) in the previous section). Then a renegotiation-proof equilibrium such as the one described under (a), featuring a combination of contributions supported by the threat of expropriation, i.e. taxes, and “voluntary” contributions is identified by the conditions:

$$\underline{c} - \underline{\theta}(v(g) - v(g - y)) = \delta(y - \underline{c}); \quad (14)$$

$$y - \bar{\theta}(v(g) - v(g - y)) = \delta(y - \bar{c}); \quad (15)$$

$$g = \underline{n}\underline{c} + \bar{n}\bar{c}. \quad (16)$$

One can verify that for the above system of equations to admit a solution with $\bar{c} > \underline{c} \geq 0$, we must have $\bar{\theta} > (\delta/(1 - \delta))\underline{\theta}$, i.e. $\bar{\theta}$ must be sufficiently larger than $\underline{\theta}$; hence, if preference types are similar, there is no scope for such a mixed equilibrium to arise.

The alternative equilibrium where only the $n/2$ higher-preference individuals pay taxes (labeled as (b) above), will be identified by the conditions

$$y - \bar{\theta}(v(g) - v(g - y)) = \delta(y - \bar{c}'); \quad (17)$$

$$g = \bar{n}\bar{c}'. \quad (18)$$

This equilibrium can dominate or be dominated by the previous one depending on the preference parameters and on the relative number of high- and low-preference individuals. If $\underline{n}/\bar{n}$ is sufficiently small, the additional revenues that can be extracted from high-preference individuals by targeting them with stronger punishment dominate the loss of revenues from low-preference individuals, and the second type of equilibrium will therefore dominate the first. If $\underline{n}/\bar{n}$ is sufficiently large, on the other hand, then the first type of equilibrium will dominate the second.

Result 20 *In the presence of unobservable preference heterogeneity, a constrained-efficient equilibrium may feature differentiated punishments with differentiated triggers, i.e. taxes paid by everyone under penalty of expropriation and “voluntary” contributions paid by a subset of individuals under penalty of a reduction in contributions by others.*

When collective consumption involves multiple collective goods, a constrained-efficient equilibrium may feature earmarking of taxes and voluntary contributions to different goods. Suppose, for example, that preferences towards a certain collective good are concentrated within a subset of individuals. Then, it will be comparatively easier to support contributions towards that good without reliance on the threat of expropriation. A constrained-efficient equilibrium may then involve earmarking of any voluntary contributions towards this good, with tax revenues being directed to the provision of the remaining public goods.

Voluntary contributions can also be instrumental in preventing expropriation attempts from occurring in equilibrium (i.e. they can be motivated by the need to preserve property rights even when they are not themselves supported by the threat of expropriation). Going back to the earlier example, where half of individuals are haves and the other half are have-nots, suppose that only a subset of the haves value public consumption highly, while the remaining haves do not. Then, the contribution level by the haves that can be supported by the preference-anonymous punishment of collective expropriation may be too small to prevent expropriation attempts by the have-nots. In this case, then, it is possible for a constrained-efficient equilibrium to feature voluntary contributions by the high-income, high-preference types that are above the tax required from high-income individuals of all preference types *under threat of* expropriation, and whose effect is *to prevent expropriation* by low income individuals. Such additional contributions can only be significant in size if the group making them is sufficiently restricted (because as we noted in Section 2, contributions supported by renegotiation proof strategies that do not rely on expropriation are limited in large groups); i.e. they must be made by a comparatively small “elite” of higher-income individuals. These additional contributions provide a collective service to the members of the elite in that they secure property rights for that elite. Compliance by individual elite members, however, is not elicited by the threat of expropriation by low-income individuals, but by the threat of a reduction in contributions by other elite members.

Result 21 *In the presence of unobservable preference heterogeneity, contributions above the tax from a subset of high-income individuals may be required in equilibrium to prevent expropriation attempts by low-income individuals from occurring.*

4 Other Dimensions of Taxation

4.1 Taxation and Incentives

Traditional models of coercive taxation have highlighted the incentive effects associated with taxation: all taxes other than lump-sum taxes alter individual economic incentives

and distort behaviour. Self-enforcing taxation can call into play incentive effects analogous to those that stem from coercive taxes, alongside incentive effects that stem specifically from the presence of enforceability constraints, as illustrated by the following discussion.

Suppose that individuals also use some of their time for leisure consumption, l^i , which affects utility in a quasilinear fashion, i.e. $u^i(x^i, l^i, g) = x^i + \theta^i v(g) + h(l^i)$, where $h'(l^i) > 0$, $h''(l^i) < 0$. Assuming the individual does not engage in expropriation, her gross of tax income is $(1 - l^i)y^i$, where y^i can be taken to represent the individual's market productivity (wage), which may or may not be private information.

Consider first a scenario where individual productivity is publicly observable. Also assume that individuals are never tempted to expropriate others ($\alpha = 0$), and that $\min\{g, g_i^*\} = g$. Along a constrained efficient equilibrium path of play, leisure will be selected so that its marginal value to the individual equals its market productivity, y^i , i.e. a level l^{i*} such that $h'(l^{i*}) = y^i$.

Suppose that income expropriation is not feasible. In this case, the binding constraint will be the repentance constraint. By the same arguments presented in Section 2, defectors can be punished most harshly by a renegotiation-proof punishment strategy that requires them to take no leisure in the punishment period, earning a market income of y^i , and contribute their full earnings to collective consumption. Hence, starting from $l^i = 0$ and $\eta^i = 1$, an alternative continuation equilibrium featuring an equal or higher level of collective consumption and resulting in a Pareto improvement could not be sustained. This punishment strategy thus guarantees renegotiation proofness.

The repentance condition—which is the binding one in this regime—in a constrained-efficient equilibrium can be written as

$$\begin{aligned} (1 - l^{i*})y^i + h(l^{i*}) - \theta^i (v(g) - v(g - y^i)) \\ = \delta \left((1 - l^{i*})y^i + h(l^{i*}) - c^i \right). \end{aligned} \quad (19)$$

Here, the punishment cost to defectors also includes the value of the leisure forgone $h(l^{i*})$ in the punishment phase; in other words, the tax that can be levied on an individual will depend on that individual's full income rather than on her earnings alone. If we totally differentiate the above with respect to y^i , we can show that since “full income” is concave in y^i the progression of the marginal tax rate with respect to y^i will be less than with fixed leisure.

When expropriation is feasible, individual i 's tax in a constrained-efficient equilibrium will be identified again by one of three possible regimes as discussed in the preceding section. In this case, however, the level of leisure taken in the punishment phase will play the same role that η^i plays in the fixed-leisure case; thus, constrained-efficient punishment strategies will involve zero leisure in the punishment phase in regimes (i) and (iii), but may allow for a positive level of leisure in regime (ii).

In all cases, as the ratio of earnings to full income increases with y^i (since l^{i*} is decreasing in y^i), the ratio of equilibrium tax to earnings will tend to fall with earnings through this

channel, implying that taxes, when measured as a proportion of earnings, will tend to fall comparatively more. However, since earnings increase at a decreasing rate as a proportion of full income (because of convexity of h), this also implies that marginal tax rate will tend to increase with earnings, even if the average tax rate is decreasing. If, on the other hand, there exist negative taxes at the lower end of the income distribution, this pattern will imply increasing average and marginal rates, with average and marginal rates flattening out at the upper end of the income distribution.

Result 22 *If income is endogenous, average tax rates will increase comparatively less (decrease comparatively more) relative to income, but marginal tax rates will increase comparatively more (decrease comparatively less).*

The above discussion applies to a scenario where ability, y^i , is fully observable. If, however, only earned income, $(1 - l^i)y^i$, is publicly observable (or verifiable), then punishment strategies must be conditioned on income and will produce labour-leisure distortions as in conventional models. To illustrate, consider an economy with two productivity types, $\underline{y}$ and $\bar{y} > \underline{y}$, each present in numbers respectively equal to $\underline{n}$ and $\bar{n}$, and suppose that the minimum contribution levels that trigger punishment are a function of earned income. Then, the respective taxes $\underline{c}$, $\bar{c}$ for each type and the levels of leisure $\underline{l}$, $\bar{l}$ taken respectively by low- and high-productivity type individuals must satisfy the enforcement constraints that apply to each respective type, as well as a no-mimicking (self-selection) constraint requiring that the high-productivity type have no incentive to mimic low productivity types by adjusting their labour supply so as to match the low-productivity type's earnings level and pay a correspondingly lower tax.^{41,42}

In such an economy, if we were to abstract from the presence of enforcement constraints, or if these constraints are not binding, information problems could only matter in the presence of a redistributive goal: if the only reason for taxation is the funding of collective consumption, any desired level of collective consumption can be achieved with coercive taxes that are uniform across productivity types, thus ensuring self-selection.

However, if the maximum common level of self-enforcing taxation that can be supported is enough to finance the efficient level of collective consumption, g^* , enforcement constraints will come into play. Starting from a common level of taxation that is sustainable for both high- and low-productivity types, additional revenue can be raised by raising the level of taxation for high-productivity types above this common level (given that their constraint is slack at that level); a higher tax on the high-productivity type, however will cause the self-selection constraint to bind.

Any constrained-efficient contract will always feature an efficient level of leisure $\bar{l} = \bar{l}^*$, since any move from this level reduces welfare without relaxing the no-mimicking constraint nor the enforceability constraint. Furthermore, the enforceability constraint for the low-productivity type will always be binding at an optimum, since relaxing this constraint—via a decrease in $\underline{c}$ below the level which makes the low-productivity type's enforcement constraint binding—results directly in lower contributions and tightens the no-mimicking con-

straint, which then needs to be offset by a reduction in contributions by high-productivity individuals and/or by a (further) reduction in $\underline{l}$ below its efficient level $\underline{l}^*$.

It is, in principle, possible for the relevant enforcement constraints for both types to be binding at an optimum and for the no-mimicking constraint to be slack—as in the a full-information scenario; in this case contributions do not distort labour supply decisions. This, however, cannot occur if $\bar{y}$ and $\underline{y}$ are close to each other.⁴³ Since we normally think of a continuum of productivity types—of which the discrete two-types model is only an abstract but convenient representation—a scenario where the self-selection constraint is slack is not practically relevant.

Then, at an optimum where the no-mimicking constraint is binding, the relevant low-productivity enforcement constraint will be binding—and will constrain the overall level of taxation in the economy—while the high-productivity enforcement constraint will be slack, with the high-productivity tax being identified by the mimicking constraint.⁴⁴ To ensure that the no-mimicking constraint is not violated, the level of labour supply of low-productivity individuals will have to lie below the efficient level, thus making mimicking more costly for high-productivity individuals. Thus, as in a standard two-class economy redistributive tax problem, an optimal self-enforcing tax structure will entail distortions in labour supply choices at the bottom end and no distortion “at the top”. However, unlike in the case of redistributive taxation, where the level of taxation on high-productivity individuals will be selected so that at the margin the marginal revenue it generates exceeds the associated marginal efficiency cost on low-productivity individuals, in this case the optimal tax structure is determined by comparison between the marginal valuation of collective consumption and the marginal cost of the distortion associated with the no-mimicking constraint.

Result 23 *In the presence of unobservable heterogeneity in market productivity, punishment strategies in a self-enforcing, constrained-efficient equilibrium may entail distortions in labour-leisure choices even in the absence of a redistributive goal.*

If the social value of a marginal revenue increase exceeds the cost of the associated marginal distortion on labour-leisure choices, it is possible for an optimum to be determined uniquely by revenue maximization considerations. At such an optimum a further increase in the high-productivity tax would require a decrease in the low-productivity labour supply, which in turn, through the enforcement constraint, would require a decrease in low-productivity taxes and thus further tighten the no-mimicking constraint; the result could then be a decrease in revenues.⁴⁵ In this regime, even in the presence of distortions, enforcement considerations alone—which here operate through a combination of self-enforcement and self-selection constraints—determine levels of taxation. In this case, the impossibility to observe (or verify) ability makes self-enforcing contributions to collective consumption distortionary just as in conventional models of distortionary coercive taxation, but the choice of contributions levels is dictated by revenue-maximization considerations.

In the alternative case where an optimum is determined by equality between the marginal valuation of collective consumption and the marginal cost of the distortion associated

with the no-mimicking constraint, the high-productivity contribution is not solely limited by enforcement incentives; as in conventional optimal tax models, it is identified as the level of tax that strikes an optimal balance between the need for increasing collective consumption provision and the cost that this entails in terms of distortions of labour-leisure choices.⁴⁶ It should be stressed, however, that even in this regime, enforceability constraints—with respect to the low-productivity contribution only in this case—are a fundamental determinant of progressivity: in the absence of any enforceability constraints (and of redistributive concerns), requiring a common tax level from all productivity types would be efficient.

Whether enforcement constraints or the costs of labour-choice distortions are going to dictate the shape of income tax schedules will depend on the extent to which labour supply can respond. If labour supply is very elastic, then consideration of labour-choice distortions will dominate; however, if labour supply is less elastic, enforceability constraints will be binding at an optimum. On the basis of this analysis, one could then formulate, the prediction that in economies characterized by comparatively less flexible responses (because of intrinsic behavioural characteristics or, equivalently, because of less flexible market institutions) enforceability considerations will be more likely to dictate the shape of income tax schedules than in economies where labour is supplied more elastically.

4.2 Preferences for Redistribution

In the preceding discussion, we have restricted our attention to cases where individuals are ex-ante risk neutral, which translates into neutrality towards ex-post inequality. This enabled us to isolate the specific implications of enforcement constraints for the structure of taxes. Inequality aversion, if present, will play a role side-to-side with enforcement considerations in determining tax progression.

Unlike in the case of coercive taxes, where optimal tax structures result from a trade-off between efficiency and equity considerations, in the case of self-enforcing taxes, enforcement constraints may dominate equity considerations. This can be shown as follows. Suppose leisure is fixed and all individuals have the same valuation for collective consumption, but there are $n/2$ high-income individuals and $n/2$ low-income individuals. Consider a no-inequality-aversion, constrained-efficient equilibrium as that which we have described, where high-income individuals pay positive taxes, low-income individuals receive a subsidy (a negative tax) and the net revenues are directed to collective consumption. If now there is inequality aversion, one possibility would be to divert some of the net revenues towards funding of private consumption. By doing so, we reduce collective consumption and increase private consumption for the poor. However, if g is sufficiently smaller than g^* , doing so will actually lower welfare for the low-income individuals, each of whom values an extra dollar worth of collective consumption more than she values an the increase in private consumption equal to $2/n$ —the per capita subsidy increase for the poor that could be funded with one dollar worth of revenue. In this case, inequality aversion will have no additional effects on the structure of taxes: taxation will be redistributive, but its structure will be independent of inequality aversion—it will be the same in a risk-neutral (utilitarian) society as in an

infinitely risk-averse society (one which only cares for the welfare of the individual who is least well off ex post). Analogously, if g is close to g^* but high-income individuals do not value collective consumption, redistribution and enforcement will both best served by a structure that directs all sustainable payments by the rich towards financing of collective consumption.

Result 24 *The structure of self-enforcing taxes in a constrained-efficient equilibrium may be invariant to changes in the degree of inequality aversion.*

On the other hand, there will be scenarios where redistribution concerns “bite”. If, for example, only the rich value collective consumption and the poor have individually no ability to expropriate others ($\alpha = 0$), a constrained-efficient tax system in the absence of inequality aversion dictates that all revenues be directed to fund collective consumption (which is only enjoyed by the rich and which the poor may nevertheless be required to fund in part), whereas in the presence of inequality aversion, a constrained-efficient equilibrium would require that some of the revenues be directed to fund private consumption for the poor. This, however, will come at a cost in terms of enforceability, because diverting funds away from collective consumption will make it more difficult to sustain payment by the rich. There exists in this case an enforcement-equity tradeoff, whereby taxes are more difficult to enforce if the revenues they generate are directed towards redistribution.

4.3 Monitoring, Evasion, Avoidance

Self-enforcing taxes as we have characterized them so far are not incompatible with the presence of monitoring activities against tax evasion, an aspect that has been extensively studied in the literature for the case of coercive taxation. The framework we have described can be easily enriched so as to capture evasion choices and monitoring activities.

It should be stressed that the problem of enforcement we have described is quite separate from the problem of tax evasion as it is usually intended. We never think of a tax evader as someone who refuses to pay a tax which she acknowledges she is liable to pay; we would describe this comparatively uncommon behaviour as a form of civil disobedience. By tax evasion we typically refer to a situation where individuals make misrepresentations to the tax authority (e.g. underdeclaring their income) in order to reduce their tax liability.

Suppose that we augment the preceding model as follows: income is exogenous but it is imperfectly observable by others; specifically, the individual who perceives the income can perfectly observe it at no cost (and faces no ex-ante uncertainty concerning this income), whereas others can only observe income at a cost through monitoring activities.

Efficiency would then require that individuals truthfully declare their income, so as to minimize the costs of public monitoring, which is a deadweight loss. However, if monitoring never takes place, individuals would always be induced to misrepresent and pay a lower tax—a tax whose level is nevertheless consistent with the application of renegotiation proof punishment as described earlier. On the other hand, constant monitoring may be too

costly. Then, constrained efficiency may require the use of randomized monitoring, just as predicted by standard analyses. However, unlike in standard analyses of tax enforcement derived from Beckerian models of crime and punishment (Becker, 1968)—which have trouble justifying the use of limited fines in conjunction with frequent monitoring rather than reliance on extreme fines—here punishment would be automatically limited by the requirements of subgame perfection and renegotiation proofness.

Our framework also makes it possible to provide a formal characterization of tax avoidance as distinct from tax evasion. Consider a situation where income is not directly observable but is believed *ex-ante* to be perfectly correlated with an observable characteristic. Then punishment strategies can be conditioned on that characteristic. Suppose that *ex post* there is a realization that correlation was not perfect as initially thought, and some individuals exhibiting a certain publicly observable characteristic actually have a higher income than that which the punishment strategies associate with that characteristic. Behaving rationally, those individuals only pay taxes that are consistent with the punishment strategies as written (i.e. they take advantage of the discovered loophole); those punishment strategies, however, were written incorrectly, although this only becomes apparent *ex post*.

In these situations, although there is no misrepresentation involved, punishment strategies could be revised *ex post* to strike at the “avoiders”. Such punishment may range from the requirement to make up retroactively for taxes unpaid to the application of the maximum sustainable combination of expropriation and fines. However, if individuals are risk averse, the possibility of punishment strategies being revised *ex post* in this way would induce an *ex-ante* risk cost on individuals. This consideration may, in a constrained efficient equilibrium, limit the application of punishment against avoiders. This would be consistent with the observation that tax administrations often impose only limited penalties on avoidance. In contrast, in any constrained-efficient equilibrium, tax evasion would always be punished as harshly as it is viable to do so.

4.4 Uncertainty

Once we interpret group cooperation as self-enforcing, anything that happens *ex post* must be consistent with a self-enforcing (equilibrium) sequence of actions. Then, it becomes natural to think again of equilibrium selection from an *ex-ante*, veil-of-ignorance perspective (Rawls, 1971). This points back to welfare theory as a predictive rather than normative tool of tax analysis: individual and collective rationality naturally suggests equilibrium selection according to *ex-ante* optimality, and so the predicted collective choice is that which maximizes *ex-ante* welfare subject to *ex-post* sustainability constraints.

Equilibrium selection according to *ex-ante* optimality does not mean that all contingencies must be explicitly dealt with *ex ante*. In principle, under each contingency, a suitable self-enforcing continuation equilibrium can be reconstructed *ex post* and identified by all individuals on the basis of the principles that guide *ex-ante* selection; once the correct equilibrium has been identified, all individuals, anticipating that other individuals will do

the same, will select the corresponding strategies as a best response.

What about the dimensionality curse that is commonly thought of as representing a fundamental obstacle to complete state-contingent contracting? In the case of binding contracts, the problem would be that it is impossible to write down a contract detailing all that the parties are to do in all possible contingencies, including contingencies that are thought of as being very unlikely. With a self-enforcing contract, however, it is not strictly necessary to do that: once the ex-ante criterion that drives equilibrium selection (e.g. ex-ante expected utility maximization) is known to all parties, then when a certain contingency arises individuals should in principle be able to reconstruct from the ex-ante criterion which self-enforcing equilibrium strategies individuals are supposed to adopt in that state of the world.

This idea may also be consistent with an environment where beliefs about the state of the world evolve as a result of learning, i.e. where individuals update their beliefs on the basis of new observations according to Bayes' rule: for example, individuals may think a certain state to be unlikely and so select a certain self-enforcing course of action, but after observing that state occurring, they may update their beliefs and move to another self-enforcing continuation equilibrium, again identified on the basis of the commonly known ex-ante selection criterion.⁴⁷ Thus, learning can be consistent with the idea that people play according to a pre-ordered scheme.

What is more difficult to accommodate in our construction is the idea that people may not have well defined beliefs about future events, e.g. a situation where something occurs which was completely unforeseen.⁴⁸ However, this is a limitation of the standard way we think about choices under uncertainty and is not specific to collective choices—as analogous problems arise in models of individual rational choice. In principle, the non-expected utility models that have been put forward in recent theoretical literature to model decisions in the presence of ill-defined priors could be adapted to our framework.

4.5 The Role of Beliefs

If any type of informational incompleteness is present, however, beliefs can play a key role in helping support cooperation under repeated interaction. Punishment strategies that are not rational in conditions of full information may be rational on the basis of beliefs that are sequentially rational under asymmetric information. To illustrate this idea, consider the case of punishment by jail.

In the equilibria described in the previous section, reductions in contributions are punished in the same way as expropriation attempts: both are punished by expropriation. In the real world, however, more severe forms of punishment—such as jail or execution—are used to punish offenders, and infringements of property rights are typically punished more severely than failing to contribute to collective consumption is. Suppose that it is possible to jail (or execute) an individual at a cost $q \geq \alpha y^i, \forall i$. Jailing an individual fully prevents the individual from attempting expropriation but cannot be an individual one-shot best response to expropriation attempts—and therefore cannot be part of one-shot equilibrium

strategies—since its cost is larger than the cost incurred by any single individual. Furthermore, if individuals can be induced not to expropriate others by other means, jailing cannot be part of an efficient, full-information equilibrium.

Suppose, however, that there exists an individual, i' , who always attempts expropriation. We can think of such an individual as an automaton, or equivalently as an individual whose discounted payoff is highest (for whatever reason) if she attempts expropriation of others than under any other circumstance. Given that i' cannot be otherwise prevented from attempting expropriation, it will be collectively optimal for the other individuals to jail i' as long as the cost of doing so is less than the damage they collectively suffer from an expropriation attempt, i.e., as long as

$$q \leq \alpha \sum_{i \neq i'} y^i. \quad (20)$$

Jailing of i' can be supported as a collective response in a constrained-efficient equilibrium through an arrangement whereby all individuals other than i' voluntarily share the cost of jailing i' (for example, contributions towards the cost of jailing may be supported by the threat of expropriation). On the other hand, it would never be rational to use jail against noncontributors, neither as an individual nor as a collective response.

If individuals can perfectly identify those who always attempt expropriation, jail cannot be used as a plausible threat as part of renegotiation-proof punishment strategies, since unlike expropriation, punishment by jailing does not generate a gain for the punishers. Nevertheless, if expropriating “deviants” cannot be identified by others *ex ante* then jail can be used as a plausible threat in a separating, renegotiation-proof, perfect Bayesian equilibrium, in which *ex-post* beliefs are derived from *ex-ante* behaviour consistently with Bayesian updating: suppose that individuals believe that only expropriation automata attempt expropriation; on the basis of this belief, it is rational for the other players to jail expropriators *ex post*; if this is how expropriation is punished, it is rational *ex ante* for individuals to expect that all expropriation attempts will be punished by jail; if, facing this threat, individuals other than expropriation automata find it optimal not attempt expropriation, the belief that only expropriation automata attempt expropriation will be consistent with behaviour (suppose, for example, that jailed individuals experience a payoff of zero; then no individual other than expropriation automata will ever attempt expropriation). Such an equilibrium will be renegotiation proof in an expected sense, because at $t+1$, if the punishers believe a defector to be an expropriation automaton, there is no scope for reCOORDINATING to a continuation equilibrium where all individuals experience a higher expected continuation payoff.⁴⁹

For jail to be a plausible punishment against expropriation, it is only necessary for such automatic expropriation behaviour to arise with some positive, albeit small, probability. Then a constrained-efficient equilibrium may involve expropriators and noncontributors being punished differentially: expropriators face jail, noncontributors face expropriation.

Result 25 *If some individual types always attempt expropriation, and individual types are unobservable ex ante, a constrained-efficient Bayesian perfect equilibrium may involve ex-*

propriators and noncontributors being punished differentially: expropriators face jail, non-contributors face expropriation and or fines.

Here, property rights will be comparatively more robust, and it will be comparatively more likely that the possibility of expropriation will contribute to support collective consumption, rather than the latter helping to support property rights.

Result 26 *If expropriation attempts can be prevented by jail, then the possibility of income expropriation unambiguously is comparatively more likely to raise the amount of collective consumption that can be sustained in comparison with an economy where expropriation is not feasible.*

When expropriation attempts can be prevented by more severe forms of punishment such as jail, there is diminished scope for collective consumption or direct transfers to be used to prevent expropriation; furthermore, potential defections will only involve a reduction in the tax paid, implying that the second term on the left-hand side of (11) and (12) will vanish, which in turn means that taxes will be positive for all income types. Furthermore, the expressions characterizing the marginal tax rate will not involve the term $-d\Lambda^i/dy^i > 0$ even for individuals for whom $\Lambda^i > 0$ —which means is that the possibility of jailing defectors will make taxes less progressive at the lower end of the income distributions.

Result 27 *If expropriators are punished by jail, taxes will be positive for all income types and comparatively less progressive than otherwise.*

Beliefs may also be might be historically determined—as discussed by Greif, Milgrom, and Weingast (1994)—and may be self-propagating. For example, suppose that are really no deviant individuals, and that active expropriations are in fact only the result of accidents. Nevertheless, starting from a belief that the probability of a deviant occurring in the population is large in comparison with the probability of a mistake occurring, it may be a best response for those who hold such belief to punish accidental deviators by death (or jail without chance of redemption), therefore never testing the belief that deviants actually exist.⁵⁰

Diverging and self-sustaining beliefs concerning the application of punishment could also help rationalize situations where only a subset of individuals comply with a formal tax that is not generally enforced. Suppose that the cost of expropriation is individual-specific and that, as a consequence, individuals are uncertain about whether punishment through expropriation would be carried out in their case. Then a subset of people may hold the belief that they would face expropriation if they defect and rationally choose to pay the tax as a result, which also implies that their prior is never tested. Others instead may hold the belief that they will not be punished by expropriation if they fail to pay, and their belief may be confirmed when they do so.⁵¹ Different and self-sustaining beliefs about the viability of punishment may also explain differences in compliance standards across groups, producing an effect analogous in practice to that of different collective behavioural “standards”, as has been described in the social norms literature.⁵²

4.6 Finite Lives

A potential objection to the idea of repeated interaction supporting cooperation is the observation that individuals have finite lives. However, the idea of self-enforcing cooperation can also be applied to environments with overlapping generations (OLG) of finitely lived individuals (Kandori, 1992). Equilibrium strategies in OLG models involve a mechanism akin to a “retirement bonus”: individuals cooperate when young and defect when old, but their defection when old will be accommodated by others (young individuals) if the old defectors have played cooperatively when young. Applied to taxation, this is consistent with unfunded social security systems as observed in many real-world systems.

In any subgame perfect equilibrium of an economy with finitely-lived individuals, individuals will always adopt their one-shot noncooperative best response in the final period of their life. Along a “cooperative” path of play, the other (young) players will accommodate this behaviour if the old player has in turn done the same when young, but will punish an old player who has failed to accommodate others when young. Punishment will thus consist of best responding to, rather than accommodating, the action chosen by the old player. The structure of such equilibria implies that renegotiation proofness will manifest itself differently in this case. Think of an economy where individuals live for two periods and where income cannot be expropriated. If a young player defects in the first period, in the last period of the defector’s life, when the punishment is applied, there is little scope for the young and the defector, now old, to recoordinate to an alternative equilibrium: given that the old player’s behaviour is simply a best response to the young players’s actions, renegotiation proofness simply requires that the action chosen by the young be optimal for them taking the old player’s reaction into account (a “Stackelberg” equilibrium). In the case of contributions to collective consumption, it will always be unilaterally optimal for an old player in a large economy to select zero contributions, so a fine cannot be part of punishment. Punishment by the young player will simply consist of collectively reducing the level of public consumption from g to a level that is collectively rational for the young. Punishment will then be particularly effective if there exist multiple forms of collective consumption some of which are old-cohort specific (i.e. used by the old and paid by the young), so that it will be collectively rational for the young to reduce provision of these goods to zero in the punishment phase.

Result 28 *In a constrained efficient equilibrium with finitely-lived individuals, contributions will be age-dependent.*

Not only will the structure of contributions not be stationary through an individual’s life, it also be nonmonotonic with respect to age. Consider an OLG framework where individuals live for three periods (young, adult, old) and all earn the same nonexpropriable income y in all periods. In each period $n/3$ old individuals die and $n/3$ new individuals are borne, thus maintaining a constant overall population size, n . Assuming that g and n are large enough that $\theta v'(g) < 1$, any individually rational strategy will involve zero contributions by old individuals, i.e. $c^3 = 0$. Adult individuals can be induced to make a

positive contribution, $c^2 > 0$ under the threat of a reduction in contributions by others in the last period of their lives:

$$c^2 - \theta (v(g) - v(g - c^3)) \leq \delta \theta (v(g) - v(\min\{g, g_3^*\})), \quad (21)$$

where g_3^* is the level of provision that is collectively optimal for young and adult individuals. Note that it is collectively rational to apply the punishment because old individuals cannot be induced otherwise to make positive contributions. Suppose now that young individuals are to be induced to make a positive contribution, $c^1 > 0$. A reduction at t could then be followed by the requirement to pay a fine at $t + 1$ followed by a reversion to normal play at $t + 2$. Note that a reduction in contributions beyond the fine at $t + 1$ and/or at $t + 2$ by others is ruled out by the need to keep other adult individuals cooperating at c^2 during punishment (so as not to violate collective rationality). Then the maximum fine payable (satisfying the repentance constraint at $t + 2$) is simply c^2 . But then it is not possible to induce young individuals to add to total tax revenues: if c^2 is lowered, positive contributions by the young may be supported, but the associated increase in contributions from the young is less than the associated loss of revenue from the adults (because we must have $c^1 = \delta(c_{MAX}^2 - c^2)$). So, a constrained-efficient equilibrium will feature $c^1 = 0$, $c^2 > 0$, $c^3 = 0$.

With income expropriation, punishment will consist of expropriation and a reduction in collective consumption for adult defectors and expropriation and a fine for young defectors, the latter supported by a punishment consisting of expropriation and a reduction in collective consumption against young defectors who do not repent when adults. Again, this will entail contributions that are first increasing and then decreasing with age. One new issue that arises naturally in connection with expropriation in a context where lives are finite is that accumulated assets may be expropriable. Then, since the profile of asset holdings is also first increasing and then decreasing with age, punishment by expropriation would also entail higher taxes for middle-aged individuals than for young and old individuals, and may in fact produce a contributions profile that is consistent with wealth rather than income taxation. If direct conditioning of taxes on age or wealth is not feasible, this may have to be achieved indirectly by differentiated taxation of labour and capital income.

5 Self-Enforcing Tax Constitutions as Social Contracts

We have characterized a constrained-efficient equilibrium as a self-enforcing social contract that is consistent with individual rationality but is selected ex-ante according to ex-ante expected utility maximization. Thus, the idea of a self-enforcing taxation takes us naturally back to the Rawlsian/Wicksellian idea of unanimous collective choices under an ex-ante veil of ignorance, which has traditionally been the intellectual foundation of utilitarianism and welfarism in general. At the same time, it also reaffirms the traditional view of individuals as utility maximizers even when acting in groups.

This interpretation is fully consistent with the traditional contractarian view of individuals consensually agreeing to a *binding* tax constitution, where ex-ante consensus guarantees

that everyone gains from the existence of taxes.⁵³ Selection of an the equilibrium that maximizes ex-ante expected utility means that the selected course of action gives everyone a higher expected payoff than any alternative (including the alternative of an indefinite repetition of one-shot noncooperative equilibria), and implies ex-ante Wicksellian unanimity. Furthermore, the continuation punishment equilibria that are used to support constrained-efficient outcomes are themselves self-supporting, renegotiation-proof equilibria; this means that, even once the veil of ignorance is lifted, each individual, even if not necessarily better off relative to *symmetric* Nash reversion, is better off along the equilibrium path than she would be under some alternative, “worst-case”, feasible equilibrium (namely, the continuation equilibrium that would be adopted to punish her). In this sense, we can say that everyone can gain from a self-enforcing tax constitution both from an ex-ante and an ex-post perspective. Thus, the characterization of self-enforcing tax constitution that we have presented in this paper follows very much in the contractarian tradition (and in doing so is counter to some more recent literature trends)—although, as noted earlier, enforceability places much more stringent requirements on the resulting structure if taxes than simple contractarianism does.

But, unlike in the traditional contractarian view, we cannot really think of ex-ante equilibrium selection as a contractual agreement in the standard sense, given that no agreement can be binding, and no contracting is therefore involved. Rather, equilibrium selection can be interpreted here as individuals each independently acknowledging that a rational group of individuals will do best, from an ex-ante perspective, coordinating to a particular self-enforcing equilibrium. And, unlike in the standard contractarian story, ex-ante equilibrium selection must account for interim and ex-post incentives, i.e. all behaviours prescribed by the selected tax constitution must be compatible with ex-post individual rationality once the veil of ignorance is broken.

Our characterization of self-enforcing tax constitutions admittedly abstracts from organizational details of real-world government institutions but can be fully reconciled with them, and can be thought of as being complementary with analyses of government delegation structures—such as, for example Laffont (2000)—that abstract from enforcement considerations. A fully developed theory of self-supporting government institutions would also need to integrate those aspects, an undertaking which is beyond the scope of the present paper.

Nevertheless, in the characterization of collective choices as resulting from a self-enforcing contract, there are two significant problems of interpretation that we wish to address before concluding our discussion. One concerns the relationship between the idea of a self-enforcing contract and politics. The other deals with the appropriateness of sticking with the idea of selfishly motivated behaviour when building models that attempt to explain collective choices.

5.1 Politics

How can we reconcile the idea that collective choices are the result of a self-enforcing contract, selected once-and-for-all *ex ante*, with the observation that real world institutions feature formalized mechanisms of collective choice—such as voting, for example? Once more we can note that groups cannot on any outside power of enforcement that could make a political constitution permanently binding; thus, whatever procedures of collective choice are adopted in equilibrium should be interpreted as being the result of a continuously-renewed, individually and collectively rational choice, and the ensuing collective choices must in turn be independently sustainable, *i.e.* self-enforcing.

The idea of *ex-ante* selection of a self-enforcing equilibrium, however, seems to leave no room for politics. Why should individuals vote, if coordination to a certain constrained-efficient outcome can be achieved spontaneously, and if the political mandate that is conferred by the vote does not *per se* convey any additional power of enforcement?

It may be tempting to argue that political competition could be thought of, in this framework, as an *ex-post* equilibrium selection mechanism: although one self-enforcing contract may be *ex-ante* optimal, other equilibria are also sustainable; then, *ex post*, once a certain distribution of individual types is realized, the different types may use the political process as an imperfect substitute for bargaining in order to agree to the *ex-post* selection of a certain outcome. This interpretation, however, clashes directly against the self-enforcing nature of cooperation. If “disagreement” occurs, what should the fallback outcome be? Given that cooperation is self-enforcing, it can be any continuation equilibrium, including those featuring cooperation.⁵⁴ Hence, a mechanical application of bargaining ideas to a repeated interaction framework appears problematic.

Another possible interpretation of politics, related to the above, is that it is necessary to deal with contractual incompleteness. The written constitutions of modern states were devised by individuals who could not have fully anticipated all contingencies, and accordingly chose to write a mainly procedural constitution, leaving it to the political process to identify the appropriate course of action in each contingency. However, as noted earlier, a self-enforcing contract does not need to have all contingencies written down, and is indeed consistent with the notion of a formal constitution as establishing basic principles. Furthermore, as we have just noted, in a self-enforcing arrangement, there is no meaningful way of interpreting politics as an ongoing bargaining process that makes up for contractual incompleteness.

Thus, the idea that tax policies reflect a self-enforcing contract calls into question the interpretation of politics as a system of procedures that substitute for bargaining in large groups—an interpretation that is common to much of the recent political economy literature.

There is, however, a possible interpretation of political processes—vaguely related to the idea of incompleteness—that is fully consistent with the idea of a self-enforcing tax constitution, namely that political competition provides an information-pooling device which may enable individuals to coordinate to a “good” equilibrium. With incomplete informa-

tion, as time goes by individuals may each have the opportunity to make new observations about the state of the world (and particularly about their own private characteristics) obtaining new information that is private to them, so there will be a problem of pooling this new private information for the purpose of Bayesian updating, which will generate scope for people trying to misrepresent what they know in order to push the outcome in directions they favour. In other words, I will try to tell others that I have new compelling information telling me that the right thing to do, according to the principles we all concur upon, is just the thing I suggest we should do, which I also happen to like. All of this could look a lot like politics.

This idea can be illustrated as follows. Consider an economy where all individuals have the same income, y , but where individuals are of different preference types, θ^i . For the sake of simplicity, we shall assume that income expropriation is not feasible. Suppose that θ^i is private information. Furthermore, suppose that—unlike in our earlier discussion—the distribution of types is also unknown. Then coordination to the self-enforcing contract that is ex-ante constrained efficient for each realization of the distribution of types must rely on individuals truthfully revealing their type. To be more specific, assume that there are only two preference types, $\underline{\theta}$ and $\bar{\theta} > \underline{\theta}$, and there are two possible realizations of the type distribution: under distribution 1, there are $\underline{n}$ and $\bar{n}$ individuals of the low and high preference types respectively, while under distribution 2, the corresponding numbers are $\underline{n} - 1$ and $\bar{n} + 1$, and each distribution can occur with probability equal respectively to π_1 and π_2 , where $\pi_1 + \pi_2 = 1$. In other words, the uncertainty over the distribution of types is limited here to the preference type of a single individual.

Suppose that individuals truthfully reveal their type under each realization; then it will be ex-ante optimal to coordinate on an arrangement whereby the strategies played by each player are the equilibrium strategies corresponding to the self-enforcing contract that is constrained-efficient for that particular realization depending on what players reveal; these contracts will be denoted as $\mathcal{C}_w$ ($w = 1, 2$). This course of action will also be individually rational, i.e. it will be consistent with a noncooperative equilibrium: if, depending on players revealing realization w ($w = 1, 2$), all individuals except i adopt the equilibrium strategies corresponding to contract $\mathcal{C}_w$, then, as long as players' announcements can be taken as truthful, adopting the equilibrium strategy that $\mathcal{C}_w$ assigns to i will be a best response for i .⁵⁵

What remains to be seen is whether, given the strategies adopted ex post in dependence of the other players' initial announcements, players will find it individually rational to reveal their true type truthfully, i.e. whether it will be a best response for each of them to announce their true type when all other players also announce their true type, and when the type distribution so revealed leads to the adoption of the self-enforcing contract that is ex-ante constrained-efficient for that realization.

It can be shown that, in this case, truthful revelation can generally not be secured. Consider the choice of an individual i who is a high-preference individual, when all other individuals truthfully reveal their type. If only $\bar{n} - 1$ of the other individuals have announced that they are high-preference types, then there is no scope for misrepresentation by i , since

it is common knowledge that a realization with $\bar{n} - 1$ high-preference individuals occurs with probability zero. However, if $\bar{n}$ of the other individuals have announced that they are high-preference types, then i , being a high-preference individual, could misrepresent her type. If the high-preference individual also announces truthfully, the resulting equilibrium is $\mathcal{C}_2$, otherwise it is $\mathcal{C}_1$. In the former, $\underline{n} - 1$ players contribute a certain level $\underline{c}_2$ and $\bar{n} + 1$ players contribute $\bar{c}_2 > \underline{c}_2$; in the latter, $\underline{n}$ players contribute a certain level $\underline{c}_1$ and $\bar{n}$ players contribute $\bar{c}_1 > \underline{c}_1$. Given that all players other than i believe the other players announcements to be truthful in each case, they will behave as dictated by the corresponding strategies. Thus, if i , being a high-preference type individual, misrepresents her type as being low-preference, all other $\bar{n}$ high-preference players will contribute $\bar{c}_1$ and i will be able to contribute $\underline{c}_1$ without incurring any punishment.

Then i will have an incentive to misrepresent her type in this way if

$$\bar{c}_2 - \underline{c}_1 - \bar{\theta}(v(g_2) - v(g_1)) > 0. \quad (22)$$

Note that, since $g_2 > g_1$, we will have (by (2)) that $\bar{c}_2 < \bar{c}_1$ and $\underline{c}_2 < \underline{c}_1$. Thus, $g_2 - g_1 = \bar{n}(\bar{c}_2 - \bar{c}_1) + (\underline{n} - 1)(\underline{c}_2 - \underline{c}_1) + \bar{c}_2 - \underline{c}_1 < \bar{c}_2 - \underline{c}_1$. Then (22) will be satisfied as long as g_1 is high enough that $\bar{\theta}v'(g_1) < 1$, since in that case the unilateral gain from the reduced contribution will exceed the i 's valuation of the reduction in collective consumption.

Low-preference individuals will have no incentives to misrepresent their type when $\underline{n} - 1$ of the other individuals (truthfully) announce that they are low-preference individuals. If they do, less than the expected level of contributions will be received (since they will not follow up their announcement with payment of the higher contribution). Punishment will then have to follow, even if individuals as a whole can infer from the deviation that misrepresentation took place. Repentance will require payment of y , and expecting for this payment to be made by the individual, i , who misrepresented her type can be an equilibrium, assuming that everyone else (including future incarnations of i) keep announcing truthfully in all subsequent rounds. Since i is punished anyway, it will be a best response for her to contribute zero. Then, the level of provision at t will be $g_2 - \bar{c}_2 < g_1$, and i will then experience a net gain, from misrepresenting her type, equal to

$$\begin{aligned} (1 + \delta)\underline{c}_1 - \delta y - \bar{\theta}(v(g_1) - v(g_2 - \bar{c}_2)) &\leq \\ (1 + \delta)\underline{c}_1 - \delta y - \bar{\theta}(v(g_1) - v(g_1 - \underline{c}_1)) &\leq 0, \end{aligned} \quad (23)$$

(since $g_2 - \bar{c}_2 < g_1 - \underline{c}_1$ and given that the second inequality coincides with (1)).

There are then two possible alternatives: either the equilibrium $\mathcal{C}_1$ is adopted in all cases, or rather than coordinating to either $\mathcal{C}_1$ or $\mathcal{C}_2$ in dependence of the players' announcements, players coordinate to alternative continuation equilibria, selected in such a way that misrepresentation by high-preference individuals does not occur. Specifically, rather than associating the ex-ante constrained-efficient equilibrium $\mathcal{C}_1$ to realization 1, players can associate an alternative contract $\mathcal{C}'_1$, for which

$$\bar{c}_2 - \underline{c}'_1 - \bar{\theta}(v(g_2) - v(g'_1)) \leq 0, \quad (24)$$

thus ensuring truthful revelation by high-preference individuals. Such alternative contract will have to feature lower aggregate contributions because of lower contributions by high-preference individuals, $\bar{c}'_1$ in comparison with the maximum supportable level $\bar{c}_1$, so as to increase the loss $\bar{\theta}v(g_2) - v(g'_1)$ relative to $\bar{c}_2 - \underline{c}'_1$, as well as higher contributions by low-preference individuals, $\underline{c}'_1$ (consistently with a lower g'_1). Such a combination $(\mathcal{C}'_1, \mathcal{C}_2)$ supported by truthful revelation will constitute an incentive-compatible, ex-ante constrained-efficient arrangement.⁵⁶

Note that anonymous revelation—whereby the distribution of individuals' announcements is publicly observable, but where announcements cannot be mapped back to the individuals making them—will be sufficient to support such an arrangement, since the selection mechanism described uses announcements in an anonymous way. Indeed, nonanonymous announcements could generally not be expected to lead to truthful revelation if the associated mechanisms uses them nonanonymously, e.g., by conditioning an individual's punishment on their announcement.⁵⁷

Result 29 *When the distribution of preferences is not publicly observable, a constrained-efficient equilibrium may involve anonymous truthful revelation. In turn this may require equilibrium contributions to deviate from their corresponding full-information levels.*

How will individuals signal their type under such an arrangement? Any set of signals will do as long as their meaning is part of an accepted set of conventions. In particular, anonymous voting over alternatives can play this role, if the various alternatives over which individuals vote are each conventionally associated to a different individual type.

Is there a natural way of associating policy alternatives to types in this way? Suppose that we think of $\mathcal{C}'_1$ and $\mathcal{C}_2$ as being the two alternatives over which voting occurs. Can we naturally associate $\mathcal{C}'_1$ with low preference and $\mathcal{C}_2$ with high preference in the sense of each being the alternative favoured by one type? Not really. It is possible for both types to be better off at $\mathcal{C}_2$ in comparison with $\mathcal{C}'_1$, assuming that each is adopted in correspondence of the true realization of the preference distribution. Thus, there does not seem to be a natural mapping from policy alternatives to signals; therefore, if the only role of voting is facilitating coordination to a self-enforcing social contract through information pooling, there does not seem to be any reason to expect voting to take place directly over policy alternatives.

On the other hand, a more natural system of signals could be provided by voting over candidate types, which provides a very natural mapping from types to signals. Then, elected policymakers would simply serve as a publicly observable signal, whose characteristics and identity convey information to individuals about which particular self-enforcing equilibrium to coordinate, even though they possess no independent policy making powers (beyond the power to enact the policies they represent).⁵⁸ Nevertheless, as discussed earlier, the need to achieve information pooling through voting will generate divergence in the policies that each of these candidate types represents—the kind of tension we associated with politics. In order to secure truthful revelation, the self-enforcing policies associated with

each candidate type may have to be skewed away from the self-enforcing policies that are ex-ante constrained efficient in each realization. In our previous discussion this translated into a bigger provision gap across policy alternatives combined.

We can generalize the above construction to the case where positive probabilities π_j are attached by individuals (with common beliefs) to all possible pairs $j = (\underline{n}, \bar{n})$ such that $\underline{n} + \bar{n} = n$. As discussed above, for any of such combinations, a high-preference individual would be tempted not to reveal her type truthfully if the constrained efficient self-enforcing contract is adopted in conjunction with it. Thus, truthful revelation will require the adoption of alternative self-enforcing contracts under each realization, such that truthful revelation will be consistent with a Bayesian Nash equilibrium, i.e. it will be a best response “on average” for high-preference individuals.⁵⁹ Again, this will require broader gaps between levels of collective consumption under each realization as well as contributions by high- and low-preference types that are closer to each other.

Note, however, that in the above story it is not only the identity of the candidate elected that conveys information about the distribution of types; it is the actual distribution of votes. This means that the collective choices that will prevail in each case will not just depend on the candidate elected by majority (as it does, for example, in a citizen-candidate model of collective choices): in this interpretation, the same candidate will adopt different policies in dependence of the voting outcome. Such prediction is not inconsistent with the observation that the actual policies put in place by policymakers do reflect the degree of political support received, i.e. the strength of the political mandate matters: a policymaker elected by a narrow margin will typically tend to adopt more centrist policies expecting that more extreme policies will be met by stronger opposition (i.e. they will hit against enforceability constraints).

The above story does not rule out that politics may play an even more substantial important information-aggregation role with respect to other types of information that may be relevant to collective choices.

Suppose, for example, that the characteristics of the contract that is ex-ante efficient depend on the prevailing state of the world, and that individuals have different prior beliefs regarding the likelihood of each state of the world prevailing. Then voting could serve as an information-aggregation device (along the lines described in the literature, e.g. by Feddersen and Pesendorfer, 1996) that could make it possible for the group to converge to a (more accurate) collective prior, and, on the basis of this aggregation, to a suitable self-enforcing contract. As time unravels and new information is gathered—some of it private but much of it public—there will a need for updating the prior on the basis of this new information, which in turn may require further pooling. For example, different individuals may acquire new pieces of (private) information and they may also be uncertain about what new information others may have acquired, which means that there will no longer be a reciprocal understanding about which continuation equilibrium everyone should be coordinating to; so, a new, single, commonly observable, announcement, replacing the various different signals privately observed by each individual will required to achieve such coordination.

Voting could take place on an ongoing basis—which may be prohibitively costly—or alternatively, if the rate of arrival of new information is not too fast, it may occur at certain specified intervals, and, in the periods intervening between consecutive voting rounds, some individual or restricted group of individuals—the appointed “policymakers”—will be relied upon to update collective beliefs on the basis of information they can gather themselves. Delegating information updating to an individual or a subset individuals in this way is more likely to be preferable to information-pooling through voting when the new information can be accessed individually by everyone, albeit at a cost. So, we can think of information about individual preferences as representing one end of the spectrum, where informational asymmetry requires pooling, while the case of public but costly information would be at the other end of the spectrum.

We can then think of appointed policymakers as individuals whose identity and type itself reveals information, and who are appointed to produce a single, publicly observable coordinating signal—their announced “policy”—an announcement that has no autonomous decision force beyond its function of inducing coordination on specific self-enforcing continuation equilibrium. In this interpretation there is a natural separation between the announcement of taxes—which is delegated to policymakers—and their enforcement—which can be delegated to independent government institutions representing the group and applying punishment consistently with collective rationality as prescribe by the particular continuation equilibrium selected.

Policymakers, in turn, may be subject to the temptation to misrepresent information they gather. To induce truthful revelation by policymakers, it may be necessary to skew continuation equilibria in their favour, which will confer some informational rents to the policymaker; however, doing so may still be ex-ante preferred to incurring the costs associated with ongoing voting.⁶⁰ In addition, truthful revelation may be supported by the application of some suitable punishment, which must itself be credible and plausible (i.e., renegotiation-proof).⁶¹ This could consist of switching to a Pareto-undominated continuation equilibrium that penalizes the policymaker given her type: in other words, a policymaker who has been found to exhibit undue bias in her choice of policy announcement may have to pay for this bias by a subsequent swing in collective choices in the opposite direction, possibly even sanctioning it while in office.⁶²

The identity of the delegate and her preferences will affect her misrepresentation incentives and the information rents that have to be surrendered to her by an ex-ante constrained-efficient system of state-contingent, self-enforcing equilibria. The delegate’s type should then be selected accordingly. This means that the ex-ante efficient delegation choice will involve delegation to a policymaker whose type is least likely, on average, to be tempted to misrepresent information. Such a choice, in combination with a suitable set of state-contingent self-enforcing equilibria (inducing truthful revelation) will minimize the expected informational rents that have to be surrendered to the delegate.⁶³

The general picture that emerges from interpreting politics in this way is still consistent with the features of real-world politics, and broadly consistent with the predictions of traditional models of political competition: political candidates will publicly disagree with

each other about which policies will be appropriate, individuals will vote for candidates according to their preferences and to the candidates' preferences as they understand them, policymakers will be selected by majority voting (see Footnote 63) and, once elected, will make choices that are as close as possible to their preferred choice. What changes is the interpretation of politics: political competition is about reconciling different opinions about the true state of the world rather than about reconciling conflicting interests—which, in principle, can be taken care of *ex ante*.

5.2 Self-Enforcing Contract or Evolved Behaviour?

There is a possible alternative explanation for group cooperation that does completely away with the problem of enforcement: the idea that cooperative behaviour could be innate to individuals within groups and be the result of evolutionary selection (see Bergstrom, 2002, for a survey).⁶⁴ The arguments that are presented in the literature combine the idea of evolutionary group selection (whereby more cooperative groups reproduce more successfully as groups) with the idea of individual selection within groups (whereby defectors reproduce more within groups) and show that the tension between these two can result in equilibria where cooperative behaviour prevails (defectors do invade groups but those groups end-up self-destructing, so that eventually only cooperators remain). Similarly, evolutionary arguments have been offered to rationalize individuals' willingness to punish defectors even when punishing is not individually rational.

One key argument against the idea that cooperation can be sustained because of hard-wired, evolutionarily selected behaviour is that there are independent reasons for selfish behaviour to be evolutionarily selected in an environment where the problem of sustaining collective consumption within groups is not the *only* problem individuals face. According to Robson and Kaplan (2003), individual optimizing behaviour coupled with learning enables individuals to cope better with a changing environment and/or to adopt a hunter/gatherer survival strategy (whereby learnt knowledge about one's environment becomes essential). Moreover, we cannot expect individual rationality with respect to output production and other individual choices to be simply decoupled from individual rationality in collective consumption choices, as it may be impossible for a rational nonopportunistic to survive in a world of rational opportunists. Consider, for example, the following variant of the evolutionary models mentioned above. Suppose that individuals interact with each other within groups both according to situations that consist of either positive-sum games (giving rise to a collective choice problem) or zero-sum games (where competition is "private" in nature), and suppose that the positive-sum or zero-sum nature of the situation is not immediately obvious to individuals. Then, individuals who selectively behave cooperatively in positive-sum situations may be vulnerable to misrepresentation by other individuals, and may therefore fare comparatively worse than purely opportunistic individuals.

Another problem of interpretation arises if we try to apply ideas from the literature on the evolution of cooperation to explain tax constitutions. Much of the literature on evolution and cooperation, although rather abstract and necessarily vague in its discussion of how

its constructs map into real-world institutions, refers to some sort of meta-game whereby people choose to behave cooperatively or noncooperatively in positive-sum situations. But there is no obvious reason why Nature should restrict itself to evolving cooperation with respect to collective consumption problems; it could also evolve suitable hard-wired behaviour in all sort of economic contexts where selfish behaviour may generate collectively undesirable consequences. So, for example, Nature should also take care of selecting people's behaviour in such a way that they do not respond at all to the presence of taxes in their budget constraints and do not take advantage of any tax avoidance opportunities, or so that, as producers, they price at marginal cost independently of profit maximization considerations (e.g. even if they happen to be monopolists). Thus, if we are prepared to abandon the idea that individuals behave in a selfish, individually rational way in the case of collective consumption, we should be prepared to do so with respect to all possible forms of group interaction.

An even more compelling argument against the idea of evolutionarily selected cooperative behaviour is that, as Robson (2001) notes, opportunistic individuals are able in any event to solve coordination failures through repeated interaction, so long as lifespans are sufficiently long. Then, the group selection advantage of hard-wired cooperators vis-à-vis opportunists would vanish, and the only thing that would remain is the comparative vulnerability of hard-wired cooperators to misrepresentations by opportunist invaders: individually rational, selfish behaviour will thus be selected over hard-wired cooperation.

On the other hand, the idea of self-enforcing tax constitutions that we put forward here is not inconsistent with the idea of evolutionary selection, as long as this takes place in parallel with repeated interaction amongst long-lived individually rational agents. There are, however, ways in which evolution and repeated interaction may interface. For example, in an overlapping generations setup, initial beliefs of newborns about their own characteristics may be systematically biased relative to population characteristics in a way that facilitates the support of cooperation without contradicting rationality and learning.⁶⁵ And, as noted earlier, under incomplete information efficient outcomes may need to be supported by a specific set of rational, self-sustaining beliefs that may be the result of evolutionary selection—which is not inconsistent with Hayek's (1967) idea that social norms are a product of cultural evolution. Furthermore, as long as fitness and utility do not diverge too much (as they should not in the long run), veil-of-ignorance selection is not inconsistent with fitness maximization, although there are difficulties in the application of expected utility theory to characterize ex-ante selection.⁶⁶ Finally, evolutionary ideas could be applied to group rather than individual selection as an alternative to ex-ante utility maximization in order to rationalize equilibrium selection of efficient outcomes under repeated interaction (Boyd and Richerson, 1994).

6 Conclusion

While the traditional public finance approach has produced valuable insights, the implicit assumption underlying that approach is that governments have the power to impose any tax they choose to. In that approach enforcement is only a problem in so far as individual attempts to evade taxes must be detected and punished. In contrast with this tradition, the informal debate on taxation has long acknowledged that general compliance depends fundamentally on the general consensus of the public—i.e. it would be impossible to enforce compliance if the public did not find the tax system acceptable and were collectively non-compliant.⁶⁷ This more fundamental question of the enforceability of tax rules has effectively been assumed away by the theoretical literature.

Early contractarian philosophers appeared to recognize that laws and institutions rest on the ongoing support of the public. Rousseau (1762), for example, notes that “sovereignty, being nothing less than the exercise of the general will, can never be alienated;” and that “the clauses ... [of the social contract] ... are so determined by the nature of the act that the slightest modification would make them vain and ineffective; so that, although they have perhaps never been formally set forth, they are everywhere the same and everywhere tacitly admitted and recognized.”⁶⁸ It is only later formalizations (such as Wicksell’s) that framed it within a theory of enforceable contracts. While contract theory can be useful for exemplifying the key ideas of social contractarianism, it supposes an individually rational but irrevocable surrender of power from individuals to government, which in turn entails very limited restrictions for the structure of tax institutions. The public finance writers in that tradition thus had to look elsewhere for principles that could sharpen their models’ predictions.

Our analysis here for the first time provides a theoretical foundation to the idea that tax rules must be viable. And, in doing so, it shows that the viability problem alone can account for observations that traditional theories can only explain by invoking distributional or ethical principles, while, at the same time, casting some light on a number of questions that traditional theories have been silent on.

Notes

¹The contractarian interpretation of the state is generally attributed to Hobbes (1651) and Locke (1690), and was later popularized by Rousseau (1762); it was first applied to the free-rider problem by Wicksell (1834), and was later popularized by several influential writers (Olson, 1965; Buchanan, 1975; Brennan and Buchanan, 1980).

²A small literature, begun by McMillan (1979), has focused on the sustainability of voluntary contributions towards collective consumption under repeated interaction, but coercive taxation as such has not been examined in this light. There is a small but growing literature on the problem of enforcement in groups (Muthoo, 2003; Dal Bo, 2001; Acemoglu, 2002). While this literature does not focus on taxation, some of the ideas in that literature are related to the questions addressed here. A recent literature dubbed “Historical and Comparative Institutional Analysis”

(Greif, 1998) uses ideas from theory of repeated games to analyze and rationalize differences in institutional structures across countries and time periods. This literature mostly focuses on cooperation in markets, but a small subset of this literature focuses on the repeated interaction between state and citizens (Weingast, 1997), starting from the traditional view that the state does have some autonomous, albeit possibly limited, power over citizens.

³For now, we shall assume that income cannot be forcibly expropriated; this assumption will be relaxed later on in our discussion.

⁴Equilibrium contributions $c^{Ni} \in [0, y^i]$, are identified by the conditions:

$$\theta^i v' \left(\sum_i c^{iN} \right) - 1 \equiv \omega^i, \quad c^{iN} \omega^i \geq 0, \quad (y^i - c^{iN}) \omega^i \leq 0, \quad \forall i. \quad (25)$$

Due to the quasilinearity of preferences, the above conditions only identify a total volume of collective consumption $g^N = \sum_i c^{iN}$, but not the distribution of contributions among individuals. Quasilinearity also implies that the efficient level of collective consumption, identified by the conditions

$$\sum_i \theta^i v'(g^*) - 1 \equiv \omega, \quad g^* \omega \geq 0, \quad \left(\sum_i y^i - g^* \right) \omega \leq 0, \quad (26)$$

is independent of how provision is funded. Conditions (25) and (26) imply $g^N \leq g^*$ (with $g^N = g^*$ if and only if $g^* = 0$). In the following discussion, we shall assume $g^* > y^i, \forall i$.

⁵The idea of renegotiation proofness relies on the Pareto ranking of alternative continuation equilibria, but there exist competing versions of this construction in the literature, which differ according to how this comparison is made. In the definition of Farrell and Maskin, renegotiation is to alternative Pareto dominant continuation equilibria which must themselves be renegotiation proof equilibria. This circularity, however, requires that the set of potential alternative equilibria be exogenously defined. According to the definition of Abreu, Pearce, and Stacchetti (1993), the comparison is done with all alternative subgame perfect continuation equilibria. Here, we follow Farrell and Maskin, but focus on a large set of comparison equilibria, namely equilibria for which collective consumption is not less than in the candidate equilibrium.

⁶It may be tempting to extend this refinement by requiring that continuation equilibria be robust to objections by *any* coalition of players, as in the Coalition-Proof Nash equilibrium refinement of Bernheim, Peleg, and Whinston (1987). That refinement, however, simply guides equilibrium selection, and is not concerned with sequential rationality in dynamic games as renegotiation proofness is. Requiring continuation equilibria to be unobjectionable for any possible coalition of players (including single players) would effectively eliminate the possibility of any form of punishment. Milder extensions, however, have been suggested—the notion of a Consistent Bargaining Equilibrium proposed by Abreu, Pearce, and Stacchetti (1993).

⁷Renegotiation-proof equilibria which are not ex-ante constrained-efficient and which involve a payment $f^i < y^i$ by the defector in the punishment phase may be possible. Suppose, for example, that all individuals have the same preferences but one-half of the population has an income of zero and the other half an income y' . Denote with $g^*(n/2)$ the level of collective consumption that is ex-ante efficient for a group of $n/2$ individuals. Then an equilibrium with $g = g^*(n/2)$ and $f' < y'$ for all individuals with income y' could be a renegotiation-proof equilibrium even if there exist another renegotiation-proof equilibrium with $f' = y'$ and a level of collective consumption above $g^*(n/2)$.

(but still short of g^*). Thus, the structure of renegotiation-proof punishment strategies remains the same even when we do not invoke ex-ante utility maximization for equilibrium selection.

⁸This entails a level of provision g_i^{**} such that $\sum_{j \neq i} \theta^j v(g_i^{**}) = 1$ if $\hat{c}_i^j < y^j, j \neq i$, but can be consistent with a level of collective consumption below g_i^{**} if $\hat{c}_i^j = y^j$ for some j . This latter possibility follows directly from the nonnegativity constraint on private consumption.

⁹Formally, we have $g_i^* = y^i + \sum_{j \neq i} \arg \min_{\hat{c}_i^j, j \neq i} \{ \sum_{j \neq i} \hat{c}_i^j \mid \hat{c}_i^j \leq y^j, j \neq i, \text{ and } \nexists \hat{c}_i^j \leq y^j, j \neq i \mid y^j - \hat{c}_i^j + \theta^i v(\sum_j \hat{c}_i^j + y^i) \geq y^j - \hat{c}_i^j + \theta^i v(\sum_j c_i^j + y^i), j \neq i, \text{ and } \hat{c}_i^j \leq y^j, j \neq i \mid y^j - \hat{c}_i^j + \theta^i v(\sum_j \hat{c}_i^j + y^i) > y^j - \hat{c}_i^j + \theta^i v(\sum_j c_i^j + y^i), \exists j \}$.

¹⁰These are respectively identified by

$$\theta^i v'(g - c^i + c^{iD}) - 1 \equiv \omega^{iD}, \quad c^{iD} \omega^{iD} \geq 0, \quad (y^i - c^{iD}) \omega^{iD} \leq 0, \quad \forall i; \quad (27)$$

$$\theta^i v'(g_i - y^i + c^{iR}) - 1 \equiv \omega^{iR}, \quad c^{iR} \omega^{iR} \geq 0, \quad (y^i - c^{iR}) \omega^{iR} \leq 0, \quad \forall i. \quad (28)$$

¹¹This can be seen by focusing on the symmetric case with $y^i = y, \theta^i = \theta, \forall i$. Denote the required payment in the punishment phase as f , with $y \geq f > c$. With n large we can write (2) as $f - \theta(v(g) - v(g - f)) - \delta(f - c) \equiv \Psi \leq 0$. As noted in the text, the expression $\Psi = 0$ is increasing in c and is greater than $c - \theta(v(g) - v(g - c)) - \delta(f - c)$ (since $f > c$), implying that the maximum sustainable contribution will be identified by the condition $\Psi = 0$. Note that $\Psi = 0$ implies $(v(g) - v(g - f))/f - (1 - \delta) = \delta(c/f) > 0$, and therefore $\theta v'(g - f) - (1 - \delta) > 0$ (since $v'(g - f) > (v(g) - v(g - f))/f$). Totally differentiating this condition, we obtain

$$\frac{\partial c}{\partial f} = \frac{\theta v'(g - f) - (1 - \delta)}{n(v'(g - f) - v'(g)) + \delta} > 0. \quad (29)$$

Thus c is maximized by raising f to its upper bound, y . Note that adding more punishment after period $t + 1$ does not make punishment more effective. This can be seen by considering an alternative punishment strategy in which the increase in the contribution required by the defector at $t + 1$ is accompanied by an increase in period $t + 2$ by an amount $z \leq y^i - c^i$, with reversion to normal play starting in period $t + 3$ rather than $t + 2$. This would increase the punishment for deviations (the right-hand side of (1) by $\delta^2 z$; the right-hand side of (2), however, would decrease by $(\delta^2 - \delta)z$. Then, since (2) always becomes binding before (1) does when $z = 0$, an increase in z would either have no effect or cause (2) to bind. Furthermore, as long as δ is sufficiently close to one, concentrating punishment in a single period is constrained efficient (it makes it possible to support the highest contribution). To see this, suppose that the fine in period $t + 1$ is reduced from y to $y - z$ but a fine γz is also applied in period $t + 2$, with $\gamma > 0$. Then, as we increase z , punishment becomes more spread out and the total punishment will increase or decrease depending on the size of γ . The repentance constraint is then $y - z - \theta v(g) + \theta(g - y + z) \leq \delta(y - c) + \delta^2 \gamma z - \delta(1 + \gamma)z$. Consider next an increase in z starting from $z = 0$. The repentance constraint that refers to period $t + 2$ will be slack (the one-shot gain γz for noncompliance is less than $y - z$ and punishment is greater than that for no-repentance at $t + 1$ since we can make punishment restart). Then totally differentiating the above, we get $\partial c / \partial z = (1 - v'(nc) + \delta^2 \gamma - \delta(1 + \gamma)) / (n(v'(g - y) - v'(g)) + \delta)$: the denominator is positive, and, for δ approaching unity, the numerator approaches $-v'(nc) < 0$ for any γ , implying that two-period punishment will be outperformed by single-period punishment. There are also

other reasons for focusing on single period punishment. For example, according to the notion of Consistent Bargaining Equilibrium proposed by Abreu, Pearce, and Stacchetti (1993) as a further refinement of renegotiation proofness in symmetric games, the punished individuals could put forward a plausible objection to a punishment strategy that skews the continuation payoff against her if this continuation payoff is below that she could attain under some alternative supportable continuation equilibrium; they show that this refinement produces solutions where punishment is concentrated immediately after defections (since a longer punishment period increases payoff asymmetry).

¹²If we differentiate (2) with respect to c and y , we obtain

$$\frac{\partial c}{\partial y} = \frac{\delta - 1 + \theta v'(g - y)}{\delta - n\theta (v'(g) - v'(g - y))}. \quad (30)$$

Since $(\delta - 1)y + v(g) - v(g - y) = \delta c \geq 0$ in equilibrium, we must have $(v(g) - v(g - y))/y > (1 - \delta)$ and therefore $v'(g - y) > 1 - \delta$ (since $v'' < 0$), which implies that (30) is positive.

¹³Note that if $g_i^* \geq g$, then it must be possible and collectively rational for the nondeviators to support a level of provision $g_i = g$ during the punishment phase, provided the defector contributes a level y^i , since the punishment for deviations from their required contributions, c_i^j , is less than that for deviations from the higher contribution levels c^i . Thus, for the punishment level of collective consumption to fall below g , it must be the case that $g_i^* < g$.

¹⁴If we focus on the symmetric case with $y^i = y, \theta^i = \theta, \forall i$, and consider renegotiation-proof equilibria supporting a constrained efficient level of collective consumption $g < g^*$, the ratio $(g^* - g_i^*)/g^*$ will tend to zero as n increases. Thus, for n large, any nonnegligible constraint on the sustainable level of provision such that $(g^* - g)/g^* > 0$ will imply $g_i^* > g$ and thus $g_i = g$. This will also be the case when individuals differ in their income levels or preferences, as long as the economy is not extremely heterogeneous. For example, if only two individuals, i' and i'' , have positive income, and all other $n - 2$ individuals have zero income, then $g_{i'}$ and $g_{i''}$ will both be different from g no matter how large n is.

¹⁵If the punishment is carried out and the defector repents, the other players benefit from a reduction in the contributions required from them in period $t + 1$, while the level of collective consumption to which they have access is collectively optimal for them. This implies renegotiation proofness, given that the punishers cannot be made better off by increasing provision to a level above g_i —since this level is collectively (constrained) optimal for them—and given that there is no scope for further compensation from the defector, who is already devoting all of her income to collective consumption.

¹⁶We shall also show there that contributions supported by the type of punishment just described can be interpreted as voluntary contributions in excess of taxes.

¹⁷Totally differentiating (2) with respect to n , with $g = nc$, we obtain

$$\left(\frac{\partial g}{\partial n}\right)_{y \text{ constant}} = \frac{\delta c}{n\theta (v'(g - y) - v'(g)) + \delta} > 0. \quad (31)$$

¹⁸Letting $y = Y/n$, with Y constant, and totally differentiating (2) with respect to n , we obtain

$$\left(\frac{\partial g}{\partial n}\right)_{ny \text{ constant}} = \frac{\delta c - (\theta v'(g - y) + (1 - \delta)) y}{n\theta (v'(g - y) - v'(g)) + \delta}. \quad (32)$$

Since $(1 - \delta)y - \theta(v(g) - v(g - y)) + \delta c = 0$ at an optimum and $v'(g - y)y > \theta(v(g) - v(g - y))$, expression (32) is negative.

¹⁹Holding $Y = ny$ constant, as n approaches infinity, collective consumption approaches the level $\hat{g} > 0$ for which $\theta v'(\hat{g}) = 1 - \delta(1 - \hat{g}/Y) < 1$, or zero if $v'(0) > 1 - \delta$; whereas the one-shot noncooperative level of collective consumption is either equal to a positive level identified by the condition $\theta v'(g^N) = 1$, or zero if $v'(0) < 1$. All of this implies $\hat{g} \geq g^N$, with the inequality being strict if $v'(0) > 1 - \delta$.

²⁰If $g_i^* < g$ for some i , punishment can also entail a fall in the level of collective provision. What this implies is that, other things equal, it will be possible to support comparatively larger contributions. Furthermore, by concavity of v , for $g_i < g$ deviation from repentance will entail a comparatively smaller gain than in the case $g_i < g$. This raises the possibility that (1) rather than (2) could become the binding condition. This is more likely the smaller is g_i^* , which in turn is related, other things equal, to the extent of income inequality: the more unequal the distribution of income is, the more likely it is that the maximum supportable level of collective consumption, g , will lie above the minimum level, g_i^* , that is consistent with constrained Pareto efficiency (as previously described).

²¹Monotonicity of (2) in c^i implies that constrained optimal trigger contribution levels are the values that make conditions (2) binding for all players, i.e. the values that solve (2) as a system of equalities. Consider the differentiation of (2) with respect to c^i and y^i , holding the total level of contributions constant. Holding g constant means holding constant aggregate income as well as its distribution; this experiment thus amounts to comparing the equilibrium contribution of two individuals with different but arbitrarily close income levels. From this experiment, we obtain

$$\left(\frac{\partial c^i}{\partial y^i} \right)_{g \text{ constant}} = \frac{\delta - 1 + \theta v'(g - y^i)}{\delta} > 0, \quad \forall i. \quad (33)$$

²²We can further differentiate (33) with respect to y^i , we obtain

$$\left(\frac{\partial^2 c^i}{\partial (y^i)^2} \right)_{g \text{ constant}} = -\frac{\theta v''(g - y^i)}{\delta} > 0, \quad \forall i. \quad (34)$$

²³For example, consider an economy where $n/2$ individuals have income $y + x$ while the remaining $n/2$ individuals have income $y - x$, and consider the effect of a mean preserving income spread, represented here by an increase in x . If we totally differentiate $g = (n/2)(c(y + x) + c(y - x))$ with respect to x we obtain $dg/dx > 0$. This result is consistent with the findings of Jun-ichi, de Meza, and Myles (1999) for the case where interaction is limited to a single round.

²⁴Totally differentiating (2), we obtain

$$\left(\frac{\partial c^i}{\partial \theta^i} \right)_{g \text{ constant}} = \frac{v(g) - v(g - y)}{\delta} > 0, \quad \forall i. \quad (35)$$

²⁵The elasticity of c^i with respect to a change in θ^i is $\theta^i (v(g) - v(g - y)) / (\delta c^i)$, which can be rewritten as $(y^i / (\delta c^i)) \theta^i (v(g) - v(g - y)) / y^i$, which is less than $(y^i / (\delta c^i)) \theta^i v'(g - y^i)$, representing the elasticity of c^i with respect to a change in y^i .

²⁶This implies that the individual's market productivity is the same as her nonmarket productivity, an assumption that will be relaxed in the next section.

²⁷As producer of the income y^j , individual j can be thought of as enjoying a natural advantage in her claim, e.g., the advantage conferred by the initial physical possession of a tradeable good, which limits the expropriation that can be carried out by others. This feature is consistent with the observation that in all societies property rights have been enforced by a combination of legal protection and physical means of protection (enclosures, locks, etc.) conferring a “possession” advantage to the owner.

²⁸As before, our discussion focuses on equilibria that are both ex-ante constrained-optimal and anonymous (identical individuals play identical strategies). Asymmetric equilibria that are ex-post constrained-efficient are in principle also possible—for example, equilibria where a subset of individuals “gang up” against one individual and expropriate her income—but these equilibria violate anonymity and are ex-ante dominated.

²⁹The expropriation-based punishment strategies so described can only be credible (subgame perfect) as part of an equilibrium where no deviations occur, and where therefore potential deviations are by single individuals. Given the technological constraints associated with expropriation, a joint deviation by a subset of individuals could not actually be punished in this way.

³⁰Here we are assuming that $\gamma n'_i < \bar{\beta}^i$, which implies that the actual income produced by a defector who fails to repent never falls short of the expropriation target $\bar{\beta}^i y^i$. If this condition were to be violated, then the repentance condition would diverge from (5) in that it would have to include the term $\max\{(\bar{\beta}^i - \gamma n'_i)y^i, 0\}$. In this case, maximizing the expropriation that is carried out against a defector would have the effect of reducing the opportunity cost of engaging in expropriation during the punishment phase. Hence, the number of punishers, s^i , in an efficient, renegotiation-proof punishment would also need to satisfy the condition $\bar{\beta}^i \leq \gamma n'_i$.

³¹Requiring that the net gain experienced by expropriators as a whole be positive is consistent with transfers being available, so that the aggregate gain can be distributed so as to ensure that each individual participating in the punishment will experience a positive gain. In turn, this guarantees that punishers will not be tempted to deviate from the stated punishment strategy, if the punishment they would incur from doing so is the same as for deviations from equilibrium play (given that the latter are themselves not profitable deviations in equilibrium, and deviating from punishment would involve an additional positive loss). The availability of transfers also raises the possibility of using fines as a substitute for expropriation; furthermore, since expropriation entails resource costs, we must consider the possibility that, upon entering the punishment phase, players may be able to renegotiate to a punishment strategy where transfers are used and where all players experience a gain in comparison with expropriation. It can be shown, however, that if we restrict our attention to single period punishment, punishment by expropriation is efficient and satisfies renegotiation-proofness. If, however, we consider multiperiod punishment we arrive at a very different conclusion: it can be shown that, for δ sufficiently large, transfers make it possible to support renegotiation-proof punishment strategies that involve indefinite payment of a fine and dominate single-period expropriation. These continuation equilibria, however, mean a permanent switch to an asymmetric outcome, and can be ruled out by imposing additional refinements, such as the Consistent Bargaining Equilibrium refinement proposed by Abreu, Pearce, and Stacchetti (1993).

³²It is possible for this requirement to be met even for $\beta(s) = \alpha$. A sufficient condition is $\alpha < \gamma\rho/(\rho - \gamma)$.

³³With $\alpha\rho/\gamma < 1$, the one-shot net gain to high-income individuals is negative, implying that they will not attempt expropriation. Then, the no-deviation condition for a high-income individual

is simply $(\alpha\rho/\gamma)y'/\delta < \delta\beta(\bar{s}')y'$.

³⁴Direct earmarking of expropriated income to collective consumption (weakly) dominates a punishment strategy whereby expropriated income is distributed to the punishers, because it reduces a punishers' ability to gainfully deviate from the stated punishment strategy during punishment.

³⁵Suppose that expropriation does not take place at $t + 1$; after giving an amount ϵ to the defector, the punishers still get y^i and save the expropriation cost; the expropriator's temptation not to repent is now increased by $\beta^i y^i$ and so the tax that can be supported at $t + 1$ is reduced by $\beta^i y^i / \delta$, whose present value to nondefectors is $\beta^i y^i > \gamma \sum_{j \leq s_j, s_j \neq i} y^j$ (since by assumption expropriation is profitable). Thus, equilibrium punishment strategies involving expropriation at $t + 1$ are renegotiation-proof (even if they entail a deadweight loss).

³⁶If, for β^i at its maximum value $\bar{\beta}^i$, the second constraint remains the binding one, applying the same arguments used in Section 2 we conclude that the effect of a higher η^i on the repentance condition is ambiguous. Then, if a decrease in η^i tightens the second constraint, η^i will be set at its upper bound $\eta^i = 1$. In this case (regime (i)) the maximum sustainable level of taxation is identified by the condition

$$(1 - \beta^i)y^i - \theta^i \left(v(\min\{g, g_i^*\}) - v(\min\{g, g_i^*\} - (1 - \beta^i)y^i) \right) + \Lambda^i - \delta \left(y^i - c^i + \theta^i (v(g) - v(\min\{g, g_i^*\})) \right) \leq 0. \quad (36)$$

In a situation where the second constraint is binding and a decrease in η^i from $\eta^i = 1$ relaxes the second-constraint (and thus raises c^i), the tax will be maximized by a value η^i below unity. However, given that $v'(g - c^i) < 1$ and $v'(g - \eta^i(1 - \beta^i)y^i) < 1$ (by assumption), for the second constraint to be binding when the first one is slack it must be the case that $c^i < \eta^i(1 - \beta^i)y^i$. Thus, η^i can only decrease to the point where $\eta^i(1 - \beta^i)y^i = c^i$ and both constraints are binding. In such a regime, with $\beta^i = \bar{\beta}^i$, and $\eta^i < 1$, the overall level of payment by a defector during the punishment phase will be less than y^i ; the maximum sustainable level of taxation in this case (regime (ii)) is identified by the no-defection condition in conjunction with the condition $\eta^i(1 - \beta^i)y^i = c^i$, which gives condition

$$\max\{c^i, 0\} - \theta^i \left(v(g) - v(g - \max\{c^i, 0\}) \right) + \Lambda^i - \delta \left(\beta^i y^i + \theta^i (v(g) - v(\min\{g, g_i^*\})) \right) \leq 0. \quad (37)$$

A corner value $\eta^i = 0$ can only occur in this case if $\Lambda^i > 0$ and $c^i \leq 0$. It is not possible to have an interior solution where $c^i / ((1 - \beta^i)y^i) < \eta^i < 1$, i.e. where only the second constraint is binding. This can be shown by noting that, since the second derivative of the left-hand side of the second constraint with respect to η^i , is equal to $\theta^i v''(\min\{g, g_i^*\} - \eta^i(1 - \beta^i)y^i) < 0$, if the first derivative with respect to η^i at $\eta^i = 1$ is positive, it remains positive for all values of η^i below unity. In a situation where $\beta^i = (\beta^i)' < \bar{\beta}^i$ and both constraints are binding, if $\eta^i < 1$, it is always possible to simultaneously increase both β^i and η^i so as to relax both constraints (for example, by increasing η^y and β^i in such a way as to hold $\eta^i(1 - \beta^i)$ constant.) Hence, the tax cannot be at a maximum in a situation where η^i and β^i both lie below their respective upper bounds and both constraints

are binding. If $\eta^i = 1$ reaches its upper bound first, with $\beta^i < \bar{\beta}^i$ and both constraints binding, the overall payment of the defector in the punishment phase is y^i , and expropriation is not used to its fully feasible extent. The maximum sustainable tax in this case (regime (iii)) is identified by the condition

$$\begin{aligned} & \max\{c^i, 0\} - \theta^i \left(v(g) - v(g - \max\{c^i, 0\}) \right) + \Lambda^i \\ & - \delta \left(y^i - c^i + \theta^i (v(g) - v(\min\{g, g_i^*\})) \right) \leq 0, \end{aligned} \quad (38)$$

³⁷If there is no income nor preference inequality ($y^i = y, \theta^i = \theta, \forall i$) and n is large and β is at its upper bound in regime (i), (12) can be written as

$$\eta(1 - \beta)y - \theta(v(g) - v(g - (1 - \beta)y)) + (\alpha\rho/\gamma - 1)y \leq \delta((\beta + \eta(1 - \beta))y - g/n), \quad (39)$$

whereas the corresponding condition without expropriation is

$$y - \theta(v(g) - v(g - y)) \leq \delta(y - g/n). \quad (40)$$

For the left-hand side of (39) to be less than the left-hand side of (40) we must have

$$(\alpha\rho/\gamma - 1 - \beta)y - \theta(v(g - (1 - \beta)y) - v(g - y)) < 0, \quad (41)$$

which is always satisfied if (7) is satisfied: if all individuals have the same income, (7) gives $\delta > 1/(\beta((\alpha\rho/\gamma) - 1))$; for $\delta < 1$ this requires $\alpha\rho/\gamma - 1 - \beta < 0$. Proceeding in the same way for regimes (ii) and (iii), we find that if (7) is satisfied, the sustainable level of collective consumption is higher with income expropriation than without it.

³⁸Warfare may also be an important channel through which the linkage between property rights and collective consumption is established. If a group is subjected to an external threat of expropriation, and if individual group members are protected from outsiders by fellow groups members—at a cost to them—then individuals who fail to contribute to collective consumption can be punished by other members of the group withholding their protection, which they can do credibly and consistently with collective rationality. This suggests a possible positive relationship between warfare and collective consumption. Modelling this relationship requires a somewhat richer framework than the one we present here, and is left for future work.

³⁹Suppose first $\Lambda^i = 0$ (implying $c^i > 0$). In regime (i), where (12) is the binding condition, we get

$$\left(\frac{\partial c^i}{\partial y^i} \right)_{g \text{ constant}} = \frac{\delta - (1 - \bar{\beta}^i - y^i d\bar{\beta}^i/dy^i)(1 - \theta v'(g - (1 - \beta(s - 1))y^i))}{\delta} > 0; \quad (42)$$

$$\left(\frac{\partial^2 c^i}{\partial (y^i)^2} \right)_{g \text{ constant}} = - \frac{(1 - \bar{\beta}^i - y^i d\bar{\beta}^i/dy^i)^2 \theta v''(g - (1 - \bar{\beta}^i)y^i) - \Xi}{\delta}. \quad (43)$$

where $\Xi \equiv (d\bar{\beta}^i/dy^i - y^i d^2\bar{\beta}^i/d(y^i)^2)(1 - \theta v'(g - (1 - \beta(s - 1))y^i))$. Unless the third derivative of β with respect to s is not positive and large, Ξ will be positive, implying that (43) will be positive. Comparing (42) with the corresponding expression for a no-expropriation scenario, (33), we can

see that the presence of expropriation results in a higher marginal rate of taxation as a function of income. In regime (ii), we have

$$\left(\frac{\partial c^i}{\partial y^i}\right)_{g \text{ constant}} = \frac{\delta(\bar{\beta}^i + y^i d\bar{\beta}^i/dy^i)}{1 - \theta v'(g - c^i)} > 0. \quad (44)$$

$$\begin{aligned} \left(\frac{\partial^2 c^i}{\partial (y^i)^2}\right)_{g \text{ constant}} &= \frac{\delta(2d\bar{\beta}^i/dy^i \bar{\beta}^i + y^i d^2\bar{\beta}^i/d(y^i)^2)}{1 - \theta v'(g - c^i)} \\ &\quad - \frac{\theta v''(g - c^i)}{1 - \theta v'(g - c^i)} \left(\left(\frac{\partial c^i}{\partial y^i}\right)_{g \text{ constant}}\right)^2 > 0. \end{aligned} \quad (45)$$

In this case in a large economy where g is close to g^* (implying that $1 - \theta v'(g - c^i)$ is close to unity), the marginal rate of taxation out of income is roughly equal to $\delta\bar{\beta}^i$, i.e. equal to the discounted, collectively-rational rate of expropriation. In regime (iii), we have

$$\left(\frac{\partial c^i}{\partial y^i}\right)_{g \text{ constant}} = \frac{\delta}{1 + \delta - \theta v'(g - c^i)} > 0; \quad (46)$$

$$\left(\frac{\partial^2 c^i}{\partial (y^i)^2}\right)_{g \text{ constant}} = -\frac{\theta v''(g - c^i)}{1 + \delta - \theta v'(g - c^i)} \left(\left(\frac{\partial c^i}{\partial y^i}\right)_{g \text{ constant}}\right)^2 > 0. \quad (47)$$

The marginal tax rate is in this case independent of $\bar{\beta}^i$, and, in a large economy where g is close to g^* , it approaches $\delta/(1 + \delta)$. However, marginal rate progression is less in regime (iii) than in regime (ii); and given that the former is more likely to apply to higher income individuals, we can conclude that rate progression will taper out with income. For $\Lambda^i > 0$ and $c^i > 0$, comparative statics effects also include the expression $-d\Lambda^i/dy^i = \rho > 0$ in the numerator, implying a higher marginal tax rate. Finally, for $\Lambda^i > 0$ and $c^i < 0$, the tax disappears from the left-hand side of the no-defection constraint, which means that regime (iii) cannot occur—because a negative c^i will be strictly less than $\eta^i(1 - \beta^i)y^i$ for all values of η^i and β^i —and the negative tax required to prevent expropriation by i is always minimized in absolute value by a choice $\eta^i = 0$. The expression for the marginal rate is then

$$\left(\frac{\partial c^i}{\partial y^i}\right)_{g \text{ constant}} = \bar{\beta}^i + \frac{\rho}{\delta} > 0, \quad (48)$$

which is greater than that for a scenario with $\Lambda^i > 0$, $c^i > 0$. Since the latter scenario is more likely to apply to comparatively lower income individuals, marginal tax rates will be comparatively higher for lower income individuals who face negative taxes. Furthermore, since Λ^i is weakly decreasing in y^i , as income increases further, Λ^i may reach zero and remain there, resulting in a discrete fall in the marginal tax rate for individuals who are not potential expropriators.

⁴⁰Optimal income tax models can predict *falling* marginal rates at the top of the income distribution, but have difficulty explaining marginal rates remaining constant beyond a certain income level.

⁴¹This constraint can be expressed as follows:

$$\left[\bar{y}(1 - \bar{l}) + h(\bar{l}) - \bar{c} \right] - \left[\underline{y}(1 - \underline{l}) + h\left(1 - \underline{y}(1 - \underline{l})/\bar{y}\right) - \underline{c} \right] \equiv \Phi(\underline{c}, \bar{c}, \underline{l}, \bar{l}; \underline{y}, \bar{y}) \geq 0. \quad (49)$$

⁴²There is another possibility to consider: the high productivity type must have no incentive to mimic the earnings of the low-productivity type *and* pay no taxes at the same time, thus behaving as a low-productivity defector (in a sense evading and avoiding at the same time). However, if the enforcement constraints for the low-productivity type are satisfied, mimicking a defecting low-productivity individual would reveal the defector as a high-productivity individual, and the appropriate punishment could then be administered; hence the standard no-mimicking constraint applies in conjunction with the relevant enforcement constraint(s)—depending on whether or not expropriation is feasible and on which optimality regime the individual falls under for the high-productivity type.

⁴³For $\underline{l} = \underline{l}^*$, $\bar{y}$ approaching $\underline{y}$, and as long as the marginal rate of contributions out of earnings is positive, mimicking by the high-productivity type produces a positive gain for that type.

⁴⁴Denote with $\underline{\Omega}(\underline{c}, \underline{l}; \underline{y}) \geq 0$ and with $\bar{\Omega}(\bar{c}, \bar{l}; \bar{y}) \geq 0$ the relevant enforcement constraints respectively for low- and high-productivity type individuals. Then, such equilibria will be characterized by $\underline{\Omega}(\underline{c}, \underline{l}; \underline{y}) = 0$, $\bar{\Omega}(\bar{c}, \bar{l}; \bar{y}) > 0$, and $\Phi(\underline{c}, \bar{c}, \underline{l}, \bar{l}; \underline{y}, \bar{y}) = 0$.

⁴⁵The condition identifying a sustainable, constrained-efficient, incentive-compatible combination $(\underline{c}, \bar{c})$ in this regime (together with the binding no-mimicking constraint and the relevant enforcement constraint for the low-productivity type) is

$$\underline{n} \frac{d\underline{c}}{d\bar{c}} + \bar{n} = 0, \quad (50)$$

where

$$\frac{d\underline{c}}{d\bar{c}} = \frac{\underline{\Omega}_{\underline{l}} \Pi_{\bar{c}}}{\underline{\Omega}_{\underline{c}} \Pi_{\underline{l}} - \underline{\Omega}_{\underline{l}} \Pi_{\underline{c}}}; \quad (51)$$

$$\frac{d\underline{l}}{d\bar{c}} = \frac{-\underline{\Omega}_{\underline{c}} \Pi_{\bar{c}}}{\underline{\Omega}_{\underline{c}} \Pi_{\underline{l}} - \underline{\Omega}_{\underline{l}} \Pi_{\underline{c}}}. \quad (52)$$

Condition (50) states that total revenues are maximized (a marginal increase in $\bar{c}$ leaves total revenues unchanged).

⁴⁶The condition identifying a sustainable, constrained-efficient, incentive-compatible combination $(\underline{c}, \bar{c})$ in this regime (together with the binding no-mimicking constraint and the relevant enforcement constraint for the low-productivity type) is

$$(n\theta v'(g) - 1) \left(\underline{n} \frac{d\underline{c}}{d\bar{c}} + \bar{n} \right) - \underline{n}(\underline{y} - h'(\underline{l})) \frac{d\underline{l}}{d\bar{c}} = 0. \quad (53)$$

with

$$\underline{n} \frac{d\underline{c}}{d\bar{c}} + \bar{n} > 0. \quad (54)$$

Condition (53) states that the marginal social surplus from the additional collective consumption that results from an increase in $\bar{c}$ must equal the marginal cost of the additional distortion induced by this increase on labour supply decisions.

⁴⁷If different individuals have access to different pieces of new information, there will be an information pooling problem: if individuals do not concur about which state of the world they are in, they will also not concur about which continuation equilibrium the ex-ante selection criterion prescribes. We return to this issue below when discussing politics.

⁴⁸We are not referring here to unforeseen contingencies which people cannot name but whose consequences may be understood. For example, individuals may be aware of the fact that certain unforeseen needs may arise with some probability (e.g. because this type of things tend to happen with some regularity) and that on average these events will affect them in a certain way (for example, requiring them to lower their private consumption). This type of uncertainty can be incorporated into ex-ante expectations. We are referring to situations where individuals are unable to form well-defined assessments of the likelihood of such unforeseen events, a situation that has been associated with the notion of “Knightian” uncertainty, and is referred to by more recent literature as situations in which individuals have ill-defined priors (as, e.g., in Chen and Epstein, 2002).

⁴⁹In reality, there may be gray areas. For example, there may be several types of “crime”, some more serious than others, and some of them may be attributable to automata (as well as to rational behaviour by non-automata) while others may not. In this case, jail may be sustainable in a Bayesian perfect equilibrium only for some types of crime and not for others, implying that there will still be scope for using transfers and/or fines/expropriation in order to deal with some type of crime. For example, murder is more easily dealt with by jail because some murderers are not motivated by economic gain; this makes it possible to apply jail/death as a penalty for murder even against murderers who perhaps are motivated by economic gain and therefore would not offend again unless it was profitable for them. On the other hand, white collar crime (financial embezzlement etc.) is harder to punish by jail because it is more obviously the result of a rational economic calculation and less the result of peculiar preferences (unless we have things such as “moral costs” being borne by most nonoffenders).

⁵⁰Such equilibria satisfy the conditions for a perfect Bayesian equilibrium, but not those for a sequential equilibrium in the sense of Kreps and Wilson (1982), since they are supported by sequentially rational beliefs for nodes which are reached with probability zero in equilibrium.

⁵¹More generally, individuals may have a limited ability to directly observe the characteristics that determine punishment incentives. In this case, compliance will be based directly on beliefs about the consequence of noncompliance rather than on a rational calculation of punishment incentives. In turn this means that individuals may be induced to comply with laws and rules which in fact violate the enforceability requirement as we have described them. However, we can think of laws and rules as having to satisfy such requirement in a “long-run sense”, where belief updating through repeated observation must give rise to beliefs that are actually consistent with punishment incentives.

⁵²Nevertheless, although our analysis has focused on constrained-efficient, high-compliance equilibria, even under full information, partial-compliance equilibria are possible alongside full-compliance ones. Then, partial-compliance equilibria could be thought of as the result of some sort of equilibrium selection “failure” on the part of a certain group. Anonymity can severely limit the possibility of partial-compliance equilibria. Consider, for example, compliance equilibria supported by expropriation-based punishment strategies. Expropriation is by its nature nonanonymous, and there will exist equilibria where punishment strategies are such that only certain individuals—identified in some fashion—are punished by expropriation if they fail to

comply. However, such equilibria are ruled out under an anonymity restriction requiring that punishment strategies be nondiscriminatory. Thus, anonymity provides another mechanism on the basis of which one can rationalize selection of full-compliance equilibria (nondiscrimination is indeed held as a key principle in most institutional systems).

⁵³Arguments for the use of benefit-based taxes are also often based on this principle (see, Brennan and Buchanan, 1980).

⁵⁴Furthermore, from an ex-ante perspective, there would be little reason to coordinate to any arrangement that allows for political competition to drive the ex-post selection of self-enforcing equilibrium, since an arrangement where coordination takes place to the ex-ante optimal equilibrium is also self-enforcing (and ex-ante optimal).

⁵⁵The problem of deriving such constrained-efficient, constrained-efficient equilibria can be formally characterized as an *implementation* problem.

⁵⁶The problem just described belongs to the class of repeated games with communication and imperfect private monitoring, as described and discussed by Compte (1998).

⁵⁷Suppose, for example, that there is income expropriation and that there are also unobservable preferences, and suppose that you have a single high-preference individual. If y is sufficiently large, and unless this individual values collective consumption enough that her marginal valuation is not below unity at g , it will not be possible to sustain higher contribution by her in comparison with low-preference individuals, whether or not the distribution of types is known (i.e. with or without truthful, anonymous revelation). On the other hand, if this individual were to reveal herself fully as high preference, this information could be used to support a higher contribution by her under threat of expropriation. If punishment strategies are conditioned nonanonymously on nonanonymous announcements in this way, they are incompatible with truthful revelation.

⁵⁸By the way, note that with yes/no questions then direct democracy may well will do the job: suppose that a policy is such that it is ex-ante constrained optimal to do it if the majority type is A but not if the majority type is B , and such that A types prefer to do it and B types don't; then simply looking at the majority voting choice and doing what they want is consistent with this story.

⁵⁹Such a best response is formally characterized by the condition

$$\sum_{j \neq (0, n)} \pi_j \left(\bar{c}'_{j+1} - \underline{c}'_j - \bar{\theta}(v(g_2) - v(g'_1)) \right) \leq 0, \quad (55)$$

where (abusing notation), if $j = (\underline{n}, \bar{n})$, $j + 1$ denotes the distribution $(\underline{n} - 1, \bar{n} + 1)$.

⁶⁰Buchanan discusses the related idea of *rational ignorance* on the part of voters in comparison with policymakers, describing policymakers as specialized information gatherers, and voters as individuals who rationally choose not to incur the cost of gathering information, preferring instead to defer decisions to the policymaker. As has been already mentioned, this choice may stem from the fact that pooling information on a continuous basis is too costly; or they may be even more compelling reasons for delegation. Suppose, for example, that there is a collective interest that certain information should not be revealed to “outsiders” (e.g. another competing group of individuals), and suppose that revealing this information to group members also reveals it to outsiders (confidential information). Then an ex-ante optimal arrangement may involve the (arbitrary) appointment of certain individuals who have exclusive access to this information and must be relied upon by others to send out suitable coordinating signals.

⁶¹This is not unlike the idea of a self-enforcing contract between voters and elected policymakers as described by Acemoglu (2002); the key difference being that policymakers do not have here unlimited latitude with respect to the choice of policies (they are effectively restricted to the self-enforcing policies that do not clash with people’s priors, although in the case of one-off policies, the enforceability constraint will not apply).

⁶²The information pooling problem can become particularly acute in the case of nonperiodic collective choices—such as whether or not to participate in a certain war—where, once the choice is made, there is no scope for further adjustment or change along that particular dimension of choice. Nevertheless, other dimensions of ongoing interaction remain, and punishment for misrepresentation by a delegate can be administered along those dimensions (e.g. a delegate who has misrepresented the case for participating in a war could be punished by an unfavourable switch to an alternative Pareto-undominated continuation equilibrium with respect to collective consumption choices).

⁶³This idea can be formalized as follows. Let θ summarize preferences, and let the continuation payoff for an individual of type θ in state s if the continuation equilibrium $\mathcal{C}$ is selected be denoted by $U(\mathcal{C}; s, \theta)$. Let $(\mathcal{C}_s(\theta'), \forall s) \equiv \mathcal{R}(\theta')$ be the profile self-enforcing state-contingent continuation equilibria (referring to all possible states) that are ex-ante constrained-efficient and induce truthful revelation in each possible state if the delegate is of type θ' (the profile $(\mathcal{C}_s, \forall s)$ which, given delegate θ' , maximizes the expectation $E_{s\theta}[U(\mathcal{C}_s; s, \theta)]$ subject to $U(\mathcal{C}_s; s, \theta') \geq U(\mathcal{C}_z; s, \theta')$, $\forall z \neq s$). Then, a constrained-efficient choice of delegate θ' is that which maximizes the expectation $E_{s\theta}[U(\mathcal{C}_s(\theta'); s, \theta)]$. Delegates with extreme preferences are on average going to be more likely to be tempted to misrepresent the true state of the world (since continuation equilibria in all states are selected ex-ante to benefit the average type). However, given that the relationship between individual types, misrepresentation incentives, and the ex-ante costs associated with ex-post informational rents is generally nonlinear, the ex-ante optimal delegate θ' need not necessarily coincide with the median nor with the mean type. The above story in itself does not require voting. All individuals would agree ex-ante that θ' is the best delegate. Suppose, however, that we also have information-pooling about individual preference to take care of and we do this by voting over candidates. So a state of the world is now represented by a couple (s, d) where d refers to a particular distribution of preferences in the population. Suppose then, that the vote must reveal individuals’ preferences and also result in a selection of delegate (consistently with the ex-ante selection of overall tax constitution, whereby we all agree that a particular delegate $\theta'(d)$ is best under distribution d). Also suppose that $\theta'(d)$ happens to be the median under d . Then voting for a certain candidate could be conventionally taken as an anonymous signal that the voter is of that type, and at the same time the median vote would actually be for $\theta'(d)$. (Note here that truthful revelation may require some further skewing of continuation equilibria to induce voters to announce truthfully (as discussed earlier) beyond the skewing required to induce truthful revelation by candidates; this can be readily formalized but is omitted here for the sake of brevity.) If, however, the median under d is not the optimal delegate $\theta'(d)$, then it could still be the case that the distance between $\theta'(d)$ and the median under d is independent of d —say equal to x —which means that casting a vote for a candidate θ'' could be taken as a signal that the voter is of type $\theta'' - x$ while ensuring that the median vote is for candidate $\theta'(d)$ under all possible type distributions. There will be situations, however, where x is not independent of d ; nevertheless voters could in principle compute the expectation $E_d(x(d))$ and vote for a type $\theta'' - E_d(x(d))$ if they are of type θ'' ; this will result in the election of a delegate that is of type $\theta'(d) - [x(d) - E_d(x(d))]$, which may result in a biased choice of candidate in some states. Such

bias will nevertheless be preferable to selecting a candidate independently of the distribution d (i.e. without voting), which would mean selecting a θ' , independent of d , that maximizes the ex-ante expectation of welfare taken over all possible distributions, d .

⁶⁴Other characterizations of cooperative behaviour that have been proposed in the literature in the context of a static collective consumption problem—such as the idea of “warm glow” in voluntary giving (Andreoni, 1990)—can also be broadly related to evolutionary ideas.

⁶⁵While the principle of insufficient reason is routinely invoked to characterize initial beliefs in the absence of observations, any initial beliefs can be consistent with learning and rationality. In particular there seems to be a lot of arbitrariness to initial beliefs concerning how experience on a certain instance of a category of objects is relevant for the whole category of objects, and specifically concerning how one’s individual preferences differ from those of others. Evolution could also manipulate beliefs of newborn individuals so as to generate an initial bias in their assessment of the likelihood of punishment being administered to defectors.

⁶⁶Individuals’ utility and evolutionary fitness should be consistent with each other in the long run. This is because behaviours that increase fitness should be reinforced, which will occur if they become more desirable from the point of view of individuals. In a symmetric, long-run evolutionary outcome, then, fitness and utility will coincide, and therefore, it will be meaningful to apply social choice constructs to characterize collective choices along a long-run equilibrium path: the constrained-efficient self-enforcing social contract (from the ex-ante point of view of individuals whose utility has been evolutionarily selected by the evolutionary process) is that which maximizes their ex-ante expected utility (as defined at that point in time), and this also maximizes fitness (since individuals make that determination in an evolutionarily stable equilibrium where their preferences are unchanged through time). However, in an asymmetric evolutionary outcome that requires individuals with different preferences to be present, we cannot really think of the preferences of the different individual types as being linked with each other in a way that is consistent with expected utility (i.e. consistent with axioms of choice over lotteries). Also, since the environment is not static and since, as a result, evolution never reaches a long-run equilibrium steady-state, applying the idea of equilibrium selection according to ex-ante welfare is problematic in an evolutionary context.

⁶⁷A relatively recent example of an unsustainable tax was the poll tax introduced in the UK in the 1980s. This turned out to be unenforceable, which is consistent with our predictions that viable taxation must satisfy certain progressivity requirement.

⁶⁸Such statements from Enlightenment Era writers can be taken as setting out a political manifesto; but they also be read as spelling out the individual and collective incentives that institutional arrangements must satisfy in order to be viable—i.e. institutional arrangements that are consistent with the “natural order” as they saw it.

References

- Abreu, D. (1988). “On the theory of infinitely repeated games with discounting,” *Econometrica* 56, 383-396.
- Abreu, D., D. Pearce, and E. Stacchetti (1993) “Renegotiation and symmetry in repeated games,” *Journal of Economic Theory* 60, 217-240.

- Acemoglu, D. (2002) "Why not a political Coase Theorem? Social conflict, commitment and politics," National Bureau of Economic Research Working Paper No. 9377.
- Andreoni, J. (1990) "Impure altruism and donations to public goods: a theory of warm-glow giving?" *Economic Journal* 100, 464-77.
- Becker, G. (1968) "Crime and punishment: an economic approach," *Journal of Political Economy* 76, 169-217.
- Bergstrom, T. (2002) "Evolution of social behavior: individual and group selection," *Journal of Economic Perspectives* 16, 67-88.
- Bergstrom, T., L. Blume, and H. Varian (1986) "On the private provision of public goods," *Journal of Public Economics* 29, 25-49.
- Bernheim, B., Peleg, B., and Whinston, M. (1987) "Coalition-proof Nash equilibria: concepts," *Journal of Economic Theory* 42, 1-12.
- Boyd, R., and P. Richerson (1994) "The evolution of norms: an anthropological view," *Journal of Institutional and Theoretical Economics* 150, 72-87.
- Brennan, G., and J. Buchanan (1980) *The power to tax: analytical foundations of a fiscal constitution*, Cambridge University Press, Cambridge.
- Buchanan, J. (1975) *The limits to liberty*. Chicago: University of Chicago Press.
- Compte, O. (1998) "Communication in repeated games with imperfect private monitoring," *Econometrica* 66, 597-626.
- Dal Bo, P. (2001) "Social norms, cooperation and inequality," UCLA Economics Working Paper No. 802.
- Chen, Z., and L. Epstein (2002) "Ambiguity, risk, and asset returns in continuous time," *Econometrica* 70, 1403-1443.
- Farrell, J., and E. Maskin (1989) "Renegotiation in repeated games," *Games and Economic Behavior* 1, 327-60.
- Feddersen, T., and W. Pesendorfer (1996) "The swing voter's curse," *American Economic Review* 86, 408-424.
- Greif, A. (1994) "Cultural beliefs and the organization of society: a historical and theoretical reflection on collectivist and individualist societies," *Journal of Political Economy* 102, 912-950.
- Greif, A. (1998) "Historical and comparative institutional analysis," *American Economic Review* 88, 80-84.
- Greif, A., P. Milgrom, and B. Weingast (1994) "Coordination, commitment, and enforcement: the case of the merchant guild," *Journal of Political Economy* 102, 745-776.
- Hayek, F. (1967) *Studies in philosophy, politics and economics*. Chicago: University of Chicago Press.
- Hobbes, T. (1651) *Leviathan*, edited by C. B. Macpherson. London: Penguin Classics, 1986.

- Jun-ichi, I., D. de Meza, and G. Myles (1999) "The private provision of public goods," *International Journal of Development Planning Literature* 14, 327-346.
- Kandori, M. (1992) "Repeated games played by overlapping generations of players," *Review of Economic Studies* 59, 81-92.
- Kreps, D., and R. Wilson (1982) "Reputation and imperfect information," *Journal of Economic Theory* 27, 253-79.
- Laffont, J.-J. (2000) *Incentives and political economy* Oxford and New York: Oxford University Press.
- Locke, J. (1690) *Two treatises on government*. Cambridge: Cambridge University Press, 2001.
- McMillan, J. (1979) "Individual incentives in the supply of public inputs," *Journal of Public Economics* 12, 87-98.
- Mirrlees, J. (1971) "An exploration of the theory of optimum income taxation," *Review of Economic Studies* 38, 175-208.
- Muthoo, A. (2003) "A model of the origins of basic property rights," University of Essex EconPaper No. 546.
- North, D. (1990) "Institutions, institutional change and economic performance," Cambridge: Cambridge University Press.
- Olson, M. (1965) *The logic of collective action: public goods and the theory of groups*. Cambridge, MA: Harvard University Press.
- Rawls, J. (1971) *A theory of justice*. Cambridge, Massachusetts: Belknap Press of Harvard University Press.
- Robson, A. (2001) "The biological basis of economic behavior," *Journal of Economic Literature* 39, 11-33.
- Robson, A., and H. Kaplan (2003) "The evolution of human life expectancy and intelligence in hunter-gatherer economies," *American Economic Review* 93, 150-69.
- Rousseau, J.-J. (1762) *The social contract* translated by M. Cranston. London: Penguin Classics, 1968.
- Van Damme, E. (1989) "Renegotiation-proof equilibria in the repeated Prisoners' Dilemma," *Journal of Economic Theory* 47, 206-17.
- Weingast, B. (1997) "The political foundations of democracy and the rule of law," *American Political Science Review* 91, 245-263.
- Wicksell, K. (1934) *Lectures on political economy*. London: Routledge, 2003.