Working Papers

Social Planning and Economic Coercion

Beat Hintermann
Thomas F. Rutherford

Social Planning and Economic Coercion

Abstract

We develop a theory of social planning with a concern for economic coercion, which we define as the difference between consumers' actual utility, and the "counterfactual" utility they expect to obtain if they were able to set policy themselves. Reasons to limit economic coercion include protecting minorities, preventing disenfranchised groups from engaging in socially costly behavior, or political economy considerations. As long as consumers are fully rational, limiting coercion is equivalent to placing more welfare weight on coerced consumers at the expense of others. If, however, consumers are not fully rational and/or informed, counterfactual utility becomes endogenous to current policy, and the welfare loss associated with limiting coercion increases. We set up a numerical version of our model and find that the error-related welfare loss can be substantial.

Beat Hintermann
Department of Economics
University of Basel
4002 Basel / Switzerland
b.hintermann@unibas.ch

Thomas F. Rutherford
Department of Economics
University of Wisconsin at Madison
Madison WI 53715 / USA
rutherford@aae.wisc.edu

October 16, 2014

1 Introduction

There are many forms of coercion, including physical, psychological, legal, religious, sexual.[1] In an economic context, coercion occurs wherever individuals are subject to decisions which directly affect their utility, but over which they have no control. Citizens of a nation-state are coerced into accepting, for example, a particular level of a public good or a tax schedule that will generally differ from what they would prefer.

Even though members of a society are coerced by group rules, joining the group may convey greater utility for an individual than remaining outside, mainly due to coercion of others e.g. in the context of enforcing contracts (Baumol, 2004). Subjection to economic coercion by joining a group can therefore be entirely voluntary; being member of the group while not being subjected to its laws would of course be even better from an individual's perspective, but this is assumed to be infeasible in a non-dictatorial society.

Gamson (1961) and Skarpedas (1992) discuss economic coercion during group formation, and Wallis (2014) and Skaperdas (2014) examine the role of coercion in emerging societies without a pre-existing rule of law.[2] In this paper, we focus on the situation where a group has already been formed and a democratic society is in place. Unlike Ledyard (2014), we assume that individuals face a high cost when leaving the society due to attachment to home, migration costs etc., such that instead of moving away, they are assumed to express their discontent by other means.

In a democratic society, economic coercion is a natural consequence of applying one set of rules to a heterogeneous population. With the exception of special cases where government policy is designed to promote the interest of a particular individual or group, all households will be economically coerced to some extent, but the degree of coercion generally varies over the population. Coercion is perhaps most obvious in the context of redistribution, where net payers are coerced at the benefit of net receivers. On the other hand, according to our definition for economic coercion, a failure to redistribute leads to a high degree of coercion for the poor. It can therefore not be the goal of government policy to eliminate coercion altogether, but there may be reasons to set limits to the amount

---

[1]Coercion is the "act, process or power of coercing", with the definition of the relevant verb given by Merriam-Webster as 1.) to restrain or dominate by force, 2.) to compel to an act or choice, or 3.) to achieve by force or threat.

[2]Many of the papers cited here are contributions to an edited volume by Martinez-Vazquez and Winer (2014), which is the result of a workshop about economic coercion held at Stone Mountain, Atlanta, in October 2010. The present paper was also presented at that workshop but does not form part of the book.

of coercion any particular individual or group may be subjected to. For example, the government may want to limit the extent to which policy choices alienate segments of the population in order to prevent them from organizing labor strikes or mass protests, or engaging in illegal activities such as refusal to pay taxes, vandalism or terrorism. Even though social welfare will be reduced by limiting coercion to some groups, the reduction in collateral damages (which in our framework are not captured by the welfare function) may more than make up for the welfare loss. Economic coercion could also be investigated in a political economy framework, where a political candidate or a ruling party may want to limit coercion to some voter groups (presumably swing voters) in order to maximize their vote share. In this paper, however, we examine the effects of limiting coercion using a social planning model.

The economic subfield where coercion within an existing group has traditionally been acknowledged is the normative theory of public choice, which is concerned with the definition of optimal decision rules for policy. This literature goes back to the work of Wicksell (1896) and Lindahl (1919), who proposed approximate unanimity as a direct consequence to the desire of minimizing coercion. Buchanan and Tullock (1962) derive unanimity in the framing of a constitution as an efficiency condition, but without providing a formal definition of economic coercion.

As discussed by Congleton (2014), defining coercion requires reference to some counterfactual or baseline. This the situation that would occur in the absence of the act or proposition in the context of which coercion is examined. For example, engaging in a binding sales contract means that the seller has to turn over the good and the buyer has to release money (or some other good, in a barter exchange); the counterfactual in this case is the absence of the contract. Naturally, this baseline could itself be associated with coercion because there are always determinants to wellbeing that are exogenous to a person. In this sense, coercion is most usefully understood as differential coercion between some counterfactual and a specific proposal.

We use as a starting point the framework developed by Winer et al. (2014), who to the best of our knowledge are the first to introduce an explicit concern with coercion into a mainstream social planning context. We adopt their working definition of coercion as the difference between the hypothetical utility a consumer thinks he should achieve, and the counterfactual utility he obtains conditional on actual policy. Whereas Winer

et al. (2014) assume that consumers arrive at their counterfactual by quantity-adjusting the level of the public good at a given tax price, in our model consumers compute their counterfactual utility by setting all policy variables. In other words, the "proposal" we have in mind is the government setting tax rates and using the revenue to supply public goods, whereas the baseline is defined by consumers making these decisions. Our paper further differs from that by Winer et al. (2014) in that the social planner limits the extent of allowable coercion to specific consumers or consumer groups, rather than aggregate coercion over the entire society (which implies some tradeoff between the coercion of different consumers), and that we introduce the coercion constraint as an inequality (i.e., an upper limit of allowable coercion) rather than a strict equality. We therefore exclude socially inferior equilibria where social welfare can be increased while simultaneously decreasing the level of coercion.

Our paper is linked to the emerging literature of behavioral public economics, which aims to incorporate systematic deviations between observed behavior on the one hand, and behavior that is consistent with neoclassical principles such as utility maximization by fully informed and rational, forward-looking agents on the other.[3] In essence, we assume that consumers compare their actual situation with what they would prefer the world to be like, and that the government has an interest to limit this difference for particular consumer groups. Whereas both assumptions are not traditional in economics, we believe that they capture an important aspect of public sentiment towards government policy. In addition, we postulate that consumers make systematic mistakes when computing their counterfactual utility, and we incorporate this assumption directly into a framework of social decision making. Considering the complex nature of modern economies, the "anomaly" of imperfectly computed counterfactual utilities based on hypothetical policy equilibrium seems quite natural. In a standard neoclassical context, consumers need not compute any counterfactual equilibrium, so the issue of whether they can do this correctly does not arise. In our context, consumers' counterfactual utility is necessary to define their level of economic coercion, and the assumption of full rationality and complete information is arguably unrealistic.

We show that under full rationality and complete information, a concern with coercion can be addressed within the social welfare function by simply adding additional welfare weight to the most coerced groups. If, however, consumers make mistakes when com-

---

[3]For a review of this literature, see Bernheim and Rangel (2007).

puting their counterfactual utility, the latter becomes endogenous, and the social planner has to consider the change in counterfactual in response to a change in policy. We also show that the error in the computation of the counterfactual leads *itself* to a welfare loss, in addition to the unavoidable loss resulting from imposing a binding coercion constraint under complete rationality. Furthermore, if counterfactual errors are sufficiently large and the coercion constraint sufficiently tight, the policy outcome may even be allocationally inefficient in the sense that it is inconsistent with social welfare maximization based on a set of nonnegative welfare weights. We derive the necessary condition for a coercion constraint to result in an inefficient outcome.

In order to obtain some intuition about the likely magnitude of the welfare loss due to limiting coercion under limited rationality, we implement a numerical version of our model with two policy dimensions and three consumer groups. Consumers derive utility from private consumption, a public good, and environmental quality. The environment is adversely affected by the use of a polluting intermediate input, which is a substitute of a clean input. We posit that consumers know the correct relationship between price changes and aggregate demand levels within markets at a given equilibrium, but that they neglect general equilibrium effects in the sense that they ignore feedbacks from other markets (e.g., the change in labor supply in response to a change in the tax rate applied to the dirty intermediary good, or the change in environmental quality in response to a change in the labor tax), and that they assume the elasticity of energy demand and labor supply to be constant even though it is not. Along all other dimensions they are fully rational. Importantly, consumers correctly incorporate the technological and government budget constraint when computing their counterfactual utilities.

We find that even with this relatively mild departure from full rationality, the welfare loss that is attributable to the endogenous nature of the counterfactual is responsible for the majority of the total welfare loss that results from imposing a binding coercion constraints. This suggests that if the government wants to limit economic coercion to some consumer groups, it is important to consider consumers' view of the world, even (or especially) if this view is not fully rational. Ignoring the error and setting policy as if consumers solved the full general equilibrium model correctly would lead to an excess in economic coercion.

The next section presents our theoretical framework and the implications of limiting

economic coercion under limited rationality for welfare and efficiency. Section 3 presents our numerical version of the model and discusses the results, and Section 4 concludes.

2 Theoretical framework

We start with a basic social planning model and then explore the welfare effects of introducing a constraint on coercion.

2.1 Model

Let X_{hi} refer to the demand for (supply of) the private good (factor) $i = 1, ..., N$ by a consumer of type $h = 1, ..., H$. Consumers derive utility from private goods and factors as well as a public good G, which may be a vector of various publicly provided goods, and which may include environmental quality. Consumer h treats the level of public provision as given and maximize his or her utility by solving

$$\max_{X_{hi}} \quad U^h(X_{hi}; G) \qquad s.t. \qquad \sum_{i=1}^{N} P_i X_{hi} \leq I_h \tag{1}$$

where I_h is a fixed amount of (non-taxable) lump-sum income, and P_i is an element of the N-dimensional vector $\mathbf{P}$ referring to the consumer price of good/factor i, which may include a tax. Solving (1) gives rise to N demand and supply functions of the form $X_{hi} = X_{hi}(\mathbf{P}, I_h, G)$, and to the indirect utility function $V^h = V^h(\mathbf{P}, I_h, G)$.

The government funds the production of the public good by means of ad-valorem taxes on the private goods and factors, with good 1 serving as the numeraire:

$$P_i = p_i(1 + t_i) \text{ for } i = 2, ..., N \tag{2}$$

Net-of-tax prices p_i are equal to the marginal rate of transformation between these goods (factors) and the numeraire. Tax revenue is fully used to finance the public good:[4]

---

[4]We assume constant production costs of the public good and normalize its unit (or their units, if G is a vector) such that the cost is one. If public production technology is general rather than linear, the cost of public provision would change with a marginal change in tax rates along with private producer prices.

$$G = \sum_{i=2}^{N} p_i \, t_i \cdot \sum_{h=1}^{H} X_{hi}(\mathbf{P}, I_h, G) \tag{3}$$

Because the focus of the model is on consumers, we abstract from the production side by assuming production to be efficient. After further substituting the market-clearing condition for goods and factors, we express the production constraint as an aggregate production-possibilities frontier of the form

$$F\left(\sum_{h=1}^{H} X_{hi}(\cdot); G \right) \le 0 \tag{4}$$

The government maximizes social welfare by choosing a vector of tax rates $\mathbf{t}$ subject to constraints on the technology and the budget requirement, and to an additional constraint on the maximally allowable amount of economic coercion that a particular consumer type may experience. From (2), producer prices are a function of tax rates and exogenous parameters representing consumers' tastes and production technology, and we can therefore express indirect utility as a function of tax rates. We use the difference between actual welfare $V^h(\cdot)$ and "counterfactual" welfare $\tilde{V}^h(\cdot)$ as our definition economic coercion.[5]

After substituting the government's budget constraint, the problem can be written as one where the social planner chooses a vector of tax rates $\mathbf{t}$ to maximize social welfare, subject to the technology and the coercion constraints:

$$\max_{\mathbf{t}} \quad W = \sum_{h=1}^{H} \alpha_h \cdot V^h\big(\mathbf{P(t)}; I_h; G(\mathbf{P(t)}, \mathbf{t}, \mathbf{I})\big) \tag{5}$$

$$s.t. \quad F\left(\sum_{h=1}^{H} X_{hi}(\cdot); G(\cdot) \right) \le 0 \tag{6}$$

$$\tilde{V}^h(\cdot) - V^h(\cdot) \le \bar{K}_h \quad \forall \, h \tag{7}$$

Individuals' utility enters social welfare according to the welfare weight α_h, with $\sum_{h=1}^{H} \alpha_h = 1$, and $\bar{K}_h$ is the upper limit of economic coercion which consumer h may experience.[6] This limit is assumed to be given, i.e. determined in the political process

---

[5]The exact form of the constraint is irrelevant due to the properties of the utility function. For example, we could transform utility by $U^* = e^U$ and use a proportional coercion constraint instead.

[6]We allow the coercion constraint to differ across consumer groups in order to avoid having to trade off one consumer's coercion against another's. Note that the coercion constraint will generally be binding for one group only, such that there is no qualitative difference between a uniform and individual coercion constraints.

along with the welfare weights. The first-order necessary condition w.r.t. to tax $t_i \in \mathbf{t}$ is

$$\sum_{h=1}^{H}(\alpha_h + \kappa_h) \cdot \frac{\partial V^h(\cdot)}{\partial t_i} - \kappa_h \frac{\partial \tilde{V}^h(\cdot)}{\partial t_i} = \lambda \cdot \frac{\partial F(\cdot)}{\partial t_i} \tag{8}$$

where λ is the shadow price on the technology constraint, and κ_h are the shadow prices of the coercion constraints for each consumer type. The derivative of the indirect utility function with respect to taxes is the sum of the effect via prices, levels of demand and supply, and the public good. Denoting aggregate demand (supply) for good (factor) i over all households as $\sum_H X_{hi} \equiv X_i$ and suppressing function arguments, the marginal change in utility is given by

$$\frac{\partial V^h}{\partial t_i} = \frac{\partial V^h}{\partial P_i} p_i + \sum_{k=1}^{N} \frac{\partial V^h}{\partial P_k} \frac{\partial p_k}{\partial t_i}(1 + t_k) + \frac{\partial V^h}{\partial G} \frac{\partial G}{\partial t_i} \tag{9}$$

$$\frac{\partial G}{\partial t_i} = \frac{p_i X_i + \sum_{k=1}^{N} \frac{\partial p_k}{\partial t_i} t_k X_k + \sum_{k=1}^{N} p_k t_k \frac{\partial X_k}{\partial t_i}}{1 - \sum_{k=1}^{N} \left(\frac{\partial p_k}{\partial G} t_k X_k + \frac{\partial X_k}{\partial G} \right)} \tag{10}$$

The first term in (9) is the direct price effect on utility via expenditure on good X_{hi}. A marginal change in t_i also affects utility via changes in the producer price vector (second term) and the public good (third term). The quantity of public provision depends on t_i via the revenue generated in market i (first term in numerator of eq. 10), net of revenue changes due to changes in prices (second term) and aggregate levels of demand and supply (third term). The denominator in (10) captures feedback effects that arise if a change in public provision affects the prices and the demand and supply of private goods and factors. Note that with linear technology, all terms involving producer price changes drop out.

The total marginal effect on the production side depends on aggregate levels of demand and supply, as well as on potential feedback effects if demand (supply) of private goods (factors) depends on the level of public provision:

$$\frac{\partial F}{\partial t_i} = \sum_{k=1}^{N} \frac{\partial F}{\partial X_k} \frac{\partial X_k}{\partial t_i} + \frac{\partial F}{\partial G} \frac{\partial G}{\partial t_i} \tag{11}$$

We now turn to counterfactual utility, which describes the utility the consumer believes he obtains under his preferred policy. In the counterfactual problem, consumers choose consumption along with the vector of tax rates $\mathbf{t}^h$ (and thus the level of public provision via the budget constraint) that they prefer, subject to a technology constraint, which may

or may not be the same as the one considered by the government. For reasons we discuss in more detail below, the solution to this problem may be depend on the current policy, which we denote as $\bar{\mathbf{t}}$ and which the consumer takes as given. The Lagrangian for consumer h's problem can then be written as

$$\max_{X_{hi}, \mathbf{t}^h} \quad \tilde{\mathscr{L}}^h = U^h\big(X_{hi}; \tilde{G}(\mathbf{t}^h, \bar{\mathbf{t}}), \mathbf{I})\big) + \tilde{\mu}_h\bigg(I_h - \sum_{i=1}^{N} P_i X_{hi}\bigg)$$
$$- \tilde{\lambda}_h \tilde{F}\bigg(\sum_{k=1}^{H} \tilde{X}_{ki}(\tilde{\mathbf{P}}(\mathbf{t}^h, \bar{\mathbf{t}}), \mathbf{I}); \tilde{G}(\mathbf{t}^h, \bar{\mathbf{t}}), \mathbf{I})\bigg) \qquad (12)$$

The "tilde" over the various functions and their dependency on current policy levels $\bar{\mathbf{t}}$ imply that consumers may not solve this problem in the exact same way the social planner would. Solving (12) and substituting the derived counterfactual demands into the utility function yields an indirect counterfactual utility function that can be abstractly written as

$$\tilde{V}^h = \tilde{V}^h[\mathbf{I}; \epsilon_h(\bar{\mathbf{t}})] \qquad (13)$$

Counterfactual indirect utility depends on exogenous income of all agents, as well as on the way consumers solve the problem, which is captured by the error $\epsilon_h(\bar{\mathbf{t}})$. This error will generally depend on the policy combination $\bar{\mathbf{t}}$ at which the counterfactual is computed. The presence of such errors, and their nature, turns out to be crucial for the characteristics of the coercion-constrained equilibrium.

Consumers can make a range of different errors when solving (12). In the most basic sense, they may fail to adequately incorporate the budget and technology constraints. Even if consumers are fully informed and rational about budget balance and production possibilities, they may not foresee all direct and indirect effects on producer prices and demand levels that result from a change in tax rates. This is the type of error that we focus on in our application, where we assume that consumers do not correctly compute the general-equilibrium effects and therefore specify preferred tax levels that are in fact not optimal for them, because the resulting prices and demand levels differ from those on which they base their expected counterfactual utility.[7]

Conversely, for $\tilde{\mathbf{P}}(\mathbf{t}^h, \bar{\mathbf{t}}) = \mathbf{P}(\mathbf{t}^h)$, $\tilde{X}_{ki}(\mathbf{P}, \mathbf{I}, \bar{\mathbf{t}}) = X_{ki}(\mathbf{P}, \mathbf{I})$, $\tilde{G}(\mathbf{P}, \mathbf{t}^h, \bar{\mathbf{t}}) = G(\mathbf{P}, \mathbf{t}^h)$ and

---

[7]A different type of "error" could occur if consumers do not simply maximize their own utility (as is assumed here) but social welfare, as has been suggested by Boadway (2014). To be consistent, this type of altruism would also have to be reflected in consumers' actual utility function.

$\tilde{F}(\cdot) = F(\cdot)$, the error drops out, and counterfactual utility becomes a function only of exogenous parameters related to preferences and production technology.[8]

The importance of the endogeneity of the consumer error w.r.t. policy can be seen by returning to the government's problem. Application of the envelope theorem implies that the derivative of the counterfactual utility $\tilde{V}^h$ w.r.t. a tax t_i is given by the corresponding derivative of the counterfactual Lagrangian, evaluated at the counterfactual demand levels and tax rates chosen by consumers. Dropping income levels as arguments for convenience, this can be written as

$$\frac{\partial \tilde{V}^h\big(\epsilon_h(\mathbf{t})\big)}{\partial t_i} = \frac{\partial \tilde{\mathscr{L}}^h\big(\epsilon_h(\mathbf{t})\big)}{\partial t_i}\bigg|_{\tilde{X}_{hi}^*, \mathbf{t}^{h*}} \tag{14}$$

If consumers solve the counterfactual problem correctly,[9] it follows that $\partial \tilde{\mathscr{L}}^h / \partial t_i = \partial \tilde{\mathscr{L}}^h / \partial t_i^h = 0 \ \forall \ i$, and the last term in brackets in (8) drops out. In this case, the welfare weight for type h is simply increased to $\alpha_h + \kappa_h \geq \alpha_h$. In other words, a concern with coercion places additional welfare weight on the consumer type(s) for which the coercion constraint is binding, but all other terms related to counterfactual utility drop out of the model. Intuitively, correctly computed counterfactuals are a function only of primitives like consumer preferences, technology and availability of resources. Since these primitives are fixed, counterfactual utility becomes a constant, and like any constant drops out of the first-order conditions. The same is true if consumers compute counterfactuals which are wrong but fixed, or counterfactuals that are wrong globally but share the same local first-order conditions. In both cases, $\partial \tilde{V}^h / \partial t_i = 0 \ \forall \ i$, and the effect of coercion is simply to place additional welfare weight on the coerced consumer type.

If, however, consumers make mistakes when computing the counterfactual, and these mistakes are sensitive to marginal policy choices, then (14) is not zero, and consumers' counterfactuals are endogenous to policy. This is made notationally explicit by making the error in (13) a function of the policy $\bar{\mathbf{t}}$ based on which the counterfactual is computed. The coercion-concerned social planner needs to take the change in counterfactual utility w.r.t. a change in a policy variable into account when maximizing social welfare. Intuitively,

---

[8]With the exception of income levels, these parameters have been suppressed in the exposition, but they are implicit in the definition of preferences and the production possibilities frontier.

[9]Equivalently, consumers can make the exact same mistakes as the government. Because we are working with a social planning model, the implicit assumption is that the government makes no mistake, but this assumption could be relaxed for example in a political economy framework where both candidates and voters may make mistakes, and which could be similar in nature.

erroneously computed counterfactual utility renders the coercion constraint a moving goal post, the trajectory of which the government has to consider when solving the coercion-constrained social welfare problem. Assuming counterfactual utility to be fixed either leads to a violation of the coercion constraint (if counterfactual utility increases), or to a nonbinding coercion constraint and thus an unnecessary loss in welfare (if counterfactual utility decreases).

2.2 Implications for welfare

By construction, adding a binding coercion constraint cannot increase welfare. The magnitude of the welfare loss depends on the stringency of the coercion constraint and on the nature of consumers' counterfactuals.

Figure 1 illustrates, using a simplified setting where the government controls two policy instruments (this is the context our numerical model in Section 3): It chooses a labor tax t_L and a tax t_E on an externality-generating good, such that every point in the figure reflects a specific policy choice. There are three consumer types $\{h, i, j\}$, who differ in their tastes and endowments, and their respective utility maxima are represented by the corners of the triangular shape.[10] The borders of this shape are the contract curves between any two types defined by the tangency of indifference curves. We call the area enclosed by the contract curves the *Pareto-optimal policy space*, because any solution within this space is allocationally efficient and could be generated as a social welfare optimum based on a particular set of nonnegative welfare weights. Any point outside the Pareto-optimal policy space therefore implies a negative welfare weight for at least one consumer type, which is inconsistent with a Paretian social welfare function.

Point B in Figure 1 represents the coercion-unconstrained welfare maximum (the "bliss point"), and the circular lines around it are social indifference curves. Suppose the coercion constraint is binding for type h, but not for types i and j (it will be binding for more than one group only as a special case). This type's indifference curves are circular lines around the utility maximum at the left corner of the policy space. An example of such an indifference curve is the line marked by I_A^h; examples of the other consumer types' indifference curves are I^j and I^i.

---

[10]The focus on three homogenous types makes the analysis simpler, but it also reflects a normative concern that households should be treated equally by the government, as argued by Wallis (2014). Rather than gearing policy to individual consumers, the government only differentiates between broad classes of consumers in our setup. In our numerical model these classes are workers and two types of capital owners.

Figure 1: Coercion-constrained policy outcome

Under full rationality, type h's counterfactual utility coincides with his or her true utility maximum. Since $\partial \tilde{V}^h/\partial t = 0$, the counterfactual is independent of actual policy, making the indifference curve also coincide with the corresponding iso-coercion curve.[11] The coercion-constrained policy outcome A is located at the tangency between the social indifference curve and type h's iso-coercion curve IC_A^h (which coincides with his indifference curve I_A^h) for a coercion level of $V_{opt}^h - V_A^h = \bar{K}_h$. The welfare loss from introducing the coercion constraint is the difference between welfare at the social optimum and the coercion-constrained outcome, W(B)-W(A).

If consumers make mistakes when computing their counterfactual in a way that makes the latter endogenous with current policy, the iso-coercion curves will no longer be tangent to the iso-utility curves because even a marginal move along an indifference curve changes the counterfactual and thus the level of coercion.[12] The coercion-constrained solution is again at the tangency between the social indifference curve and type h's relevant

---

[11]A completely correct counterfactual is a sufficient but not a necessary condition. A wrong counterfactual that is fixed, or a counterfactual that is wrong globally but correct locally satisfy the condition $\partial \tilde{V}^h/\partial t = 0$ as well.

[12]The slope of the indifference curve is $\frac{dt_L}{dt_E} = -\frac{\partial V/\partial t_E}{\partial V/\partial t_L}$, whereas the slope of the iso-coercion curve is given by $\frac{dt_L}{dt_E} = -\frac{\partial \tilde{V}/\partial t_E - \partial V/\partial t_E}{\partial \tilde{V}/\partial t_L - \partial V/\partial t_L}$. The two only coincide if $\frac{\partial \tilde{V}}{\partial t_L} = \frac{\partial \tilde{V}}{\partial t_E} = 0$.

iso-coercion curve, for example $IC_{A'}$, leading to a solution at A′. The resulting welfare loss in this example is W(B)-W(A′).

The endogeneity of the counterfactual leads to a welfare loss in itself. To see this, suppose we could somehow fix the counterfactual associated with A′ in consumer type h's mind, such that $\partial \tilde{V}^h / \partial t = 0$ and the iso-coercion curve coincides with consumer h's indifference curve I_A. The resulting coercion-constrained policy outcome would be A as in the case of a correct counterfactual. The welfare loss due to the endogenous nature of the counterfactual is therefore W(A)-W(A′), whereas W(B)-W(A) can be interpreted as the unavoidable loss of introducing a coercion constraint. The move from policy A′ to policy A cannot take place because it would make consumer h adjust his counterfactual utility such that the coercion constraint is exceeded. Correcting the error embedded in the calculation of the counterfactual allows society to reach point A, which provides a greater welfare for the same level of (corrected) coercion.

In general, welfare is increased by fixing the counterfactual whenever the indifference curve of the most-coerced consumer (i.e., the consumer type for whom the coercion constraint is binding) is not tangent to the iso-coercion curve at the equilibrium, which is the case if and only if $\partial \tilde{V}^h / \partial t \neq 0$. Intuitively, the error embedded in the computation of counterfactual utility can be interpreted as an additional constraint on the welfare maximization problem, and like any additional constraint it can only reduce welfare.

2.3 Efficiency

If consumer errors are large and the coercion constraint sufficiently tight, an errouneous counterfactual can lead to a policy equilibrium that is not only distributionally, but even allocationally inefficient. An example of such an outcome is point A″ in Figure 2, located at the tangency between the iso-coercion curve $IC_{A''}$ and the corresponding social indifference curve. Starting from point A″, moving inside the shaded area would not only increase social welfare and (weakly) type h's utility, but also allocational efficiency. Unlike the equilibria at A and A′ in Figure 1, the solution at A″ could not be replicated by a set of nonnegative welfare weights, since all such solutions have to lie within the Pareto-optimal policy space. An equilibrium at a point like A″ therefore corresponds to an overall negative welfare weight for at least one agent.

The understand the nature of the negative welfare weight due to the introduction of a

Figure 2: Allocationally inefficient outcome
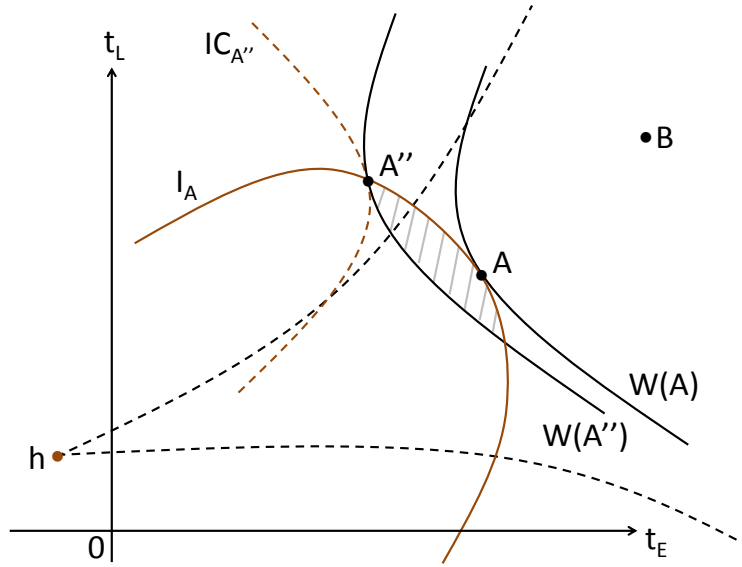
coercion constraint, we can interpret the additional terms in (8) as welfare weights that differ across policy dimensions. These weights can be positive or negative, depending on the coercion response to a change in a policy variable. Setting the LHS equal to zero and solving for α_h shows that consumer type h receives a negative welfare weight along policy dimension $t_i \in t$ if

$$\alpha_h < \kappa_h \cdot \left(\frac{\partial \tilde{V}^h / \partial t_i}{\partial V^h / \partial t_i} - 1 \right) \tag{15}$$

An inefficient result such as point A″ in Figure 2 implies that condition (15) applies to at least one policy dimension, and that the negativity of the dimension-specific welfare weight is sufficiently large to render the overall welfare weight negative. This condition can only occur if the coercion constraint for type h is binding, and if the marginal effect of t_i on counterfactual utility has the same sign and a greater magnitude than the effect on actual utility. Suppose that $\partial V^h / \partial t_i > 0$ at the equilibrium, because the utility gain via an increase in the public good more than offsets the utility loss from a higher price. If counterfactual utility, and thus coercion, increases by more than actual utility and α_h / κ_h is sufficiently low such that (15) holds, the government will consider this consumer type's preferences negatively when computing the optimal tax on X_i. The reason is that the increase in welfare due to an increase in the consumer's utility is more than offset by the increase in economic coercion.

The coercion-constrained outcome under an erroneous counterfactual can be inter-

14

preted as a type of paternalism, but one that refers to consumers' counterfactual rather than actual utility function. We refer to this situation as counterfactual paternalism: Not taking the changing nature of the counterfactual into account, consumer type h thinks he would prefer a move from A′ inside the shaded area in Figure 1, or from A″ inside the shaded area in Figure 2. The social planner does not allow such a move even though it would increase overall welfare (and possibly even efficiency), because he knows that this would result in economic coercion beyond the limit $\bar{K}_h$, once the counterfactual adjusts.

2.4 Minimizing coercion

Unless consumers have homogenous preferences and identical endowments, every policy choice will lead to some coercion since policy variables apply to everybody. A certain level of coercion is therefore unavoidable, even if the government were to single-mindedly pursue a policy of minimizing coercion. Naturally, economic coercion could be eliminated for some consumer group, but only at the price of significant coercion of others.

The minimum level of coercion that equally applies to all consumers[13] can be found be minimizing $\bar{K}_h = \bar{K} \,\forall\, h$ subject to the production possibility frontier and the coercion constraint:

$$\min_t \quad \bar{K} \qquad s.t. \quad F(\cdot) \leq 0$$
$$\tilde{V}^h(\cdot) - V^h(\cdot) \leq \bar{K} \qquad (16)$$

Figure 3 shows the relationship between the coercion constraint and social welfare. The slope of this figure is given by differentiating the government's Lagrangian w.r.t. the coercion constraint:

$$\frac{\partial W}{\partial \bar{K}} = \left.\frac{\partial \mathscr{L}}{\partial \bar{K}}\right|_{t^*} = \sum_{h=1}^{H} \kappa_h$$

The minimum attainable level of coercion is given by $\bar{K}^{min}$, which is associated with a level of social welfare of $\bar{W}$. Coercion-unconstrained social welfare is maximized at the point $(\bar{K}^0, W^0)$, where all shadow prices κ_h are zero. Increasing $\bar{K}$ beyond this point has no effect on welfare in our model (solid line) as the constraint is no longer binding.

---

[13]Note that even if all consumers are subject to the same numeric constraint, this could imply different levels of coercion if tastes and consumer errors are heterogeneous.

Figure 3: Social welfare as a function of the constraint on coercion
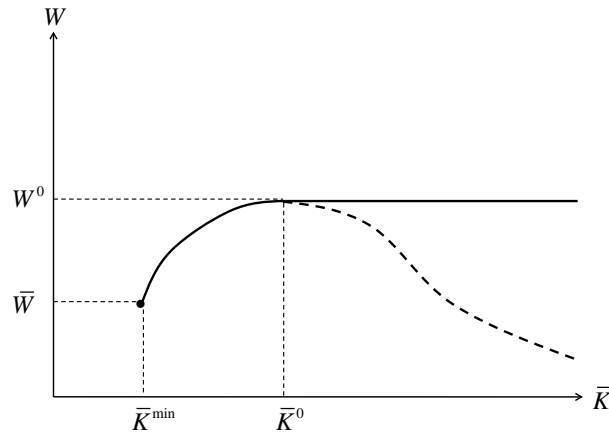
If one were to impose a coercion constraint that holds with a strict equality as in Winer et al. (2014), rather than the weak inequality used in (6), increasing $\bar{K}$ beyond $\bar{K}^0$ would decrease social welfare as indicated by the dashed line. An outcome on this line would not be socially desirable because lowering the level of coercion would both alleviate social tensions for the most coerced and increase overall welfare.

If instead the goal is to minimize coercion to one particular consumer, the solution consists in choosing the tax combination that solves the following equation as a fixed point:

$$\mathbf{t}^h_{fp} = \operatorname{argmax} \left\{ \tilde{\mathscr{L}}^h\big(\epsilon_h(\mathbf{t}^h_{fp})\big) \right\} \tag{17}$$

Naturally, if consumers make no mistake when computing their counterfactual utility, the minimum level of coercion occurs at a consumer's true utility maximum. If, however, the counterfactual is endogenous to the policy from which it is computed, the consumer may not want to move away from a policy combination even if he would in fact be better off at a different equilibrium. The reason he does not want to move there is that he mis-computes the utility associated with his true utility maximum, which makes him believe that he is better off at the fixed point. We will return to the fixed points in our numerical application.

3 Numerical application

To assess the magnitude of welfare changes from introducing a coercion constraint in the presence of consumer errors, we develop a numerical version of our model by choosing specific functional forms to represent preferences and technology. We simplify the theoretical model by restricting the number of policy instruments to a tax on labor and a tax on an intermediate good E that is associated with an aggregate environmental externality, and to three consumer types. The model presented here is the non-calibrated version, which is easier for exposition purposes. For the actual solution we transform everything into a calibrated share-form relative to a benchmark.[14]

3.1 Specification of functional forms

We represent utility by a nested constant elasticity of substitution (CES) function. In the top nest, a private consumption composite C_h and leisure ℓ_h are combined according to a Cobb-Douglas relationship. This private aggregate then produces utility along with public provision G and environmental quality Q:

$$U^h = U^h(C_h, \ell_h, G, Q) = \left[\left(\theta_h^C + \theta_h^\ell \right) \left(C_h^{\delta_h} \ell_h^{1-\delta_h} \right)^{\rho_u} + \theta_h^G G^{\rho_u} + \theta_h^Q Q^{\rho_u} \right]^{1/\rho_u} \qquad (18)$$

The parameter δ_h reflects the share in total income spent on private consumption goods within the Cobb-Douglas nest; the θ_h^k, $k = C, \ell, G, Q$, are the share parameters of the CES function with $\sum_k \theta_h^k = 1 \, \forall \, h$; and the exponent $\rho_u \equiv 1 - 1/\sigma_u$ reflects the curvature of the indifference curves, where σ_u is the elasticity of substitution between the private aggregate, public provision and environmental quality.

Consumers maximize their utility subject to the budget constraint

$$M_h = p_L L_h + I_h \leq C_h; \qquad L_h = \bar{L}_h - \ell_h \qquad (19)$$

where L_h refers to consumer h's labor supply, p_L is the (uniform) net-of-tax wage and $\bar{L}_h$ represents the time endowment in efficiency units.[15] In addition to income from labor, consumers may also receive non-taxable income I_h.

---

[14]The GAMS file used to implement and solve the model is available from the authors upon request.

[15]This allows us to use one wage rate that applies to all consumers. Consumer heterogeneity is captured by variation in $\bar{L}_h$.

Consumers' shadow value of time is given by ω_h. The comparative slackness condition

$$\omega_h \geq p_L; \quad L_h \geq 0; \quad (\omega_h - p_L) \cdot L_h = 0 \tag{20}$$

allows for the possibility of consumers not entering the labor market, if their valuation of time exceeds p_L. The consumer demand for private consumption and leisure resulting from maximizing (18) subject to (19) is

$$C_h = \frac{\theta_h^C M_h}{p_C}$$
$$\ell_h = \frac{\theta_h^L M_h}{p_L} \quad \text{if} \quad \omega_h = p_L; \qquad \ell_h = \bar{L}_h \quad \text{otherwise}$$

Production takes place in two stages. In the first stage, labor and I_h are used linearly to produce two intermediate goods:

$$\sum_h (L_h + I_h) = X + E \tag{21}$$

Good X is a clean composite, whereas E is associated with an externality that negatively affects environmental quality according to

$$Q(E) = E^{-\gamma} \qquad \gamma > 0$$

Due to the linear technology, producer prices for X and E are fixed, and we can choose quantities such that they are unity. In a second production stage, they are used to produce final output C and G according to a CES production function:

$$f(X, E) = Y \equiv G + \sum_h C_h$$
$$= \Phi \left[\theta_X X^{\rho_y} + (1 - \theta_X) E^{\rho_y} \right]^{1/\rho_y}; \qquad \rho_y = 1 - 1/\sigma_y \tag{22}$$

Here, $0 \leq \theta_X \leq 1$ is a share parameter, σ_y is the elasticity of substitution between X and E in the production of final output Y, and Φ is a constant that will take on the value of expenditure in the benchmark to express utility in money-metric terms.

The government chooses ad-valorem tax rates for labor and the dirty intermediate good, t_L and t_E, respectively, but taxes neither the consumption composite nor the clean

intermediate good X. With-tax prices for E and labor are given by $P_E = p_E(1 + t_E)$ and $P_L = p_L(1 + t_L)$, respectively. Tax revenue is used to fund the public good:

$$p_Y G = t_L p_L \sum_h L_h + t_E E \tag{23}$$

With CES technology, the price of the final output (equal to the unit cost of production, which applies to both C and G) is

$$p_Y = \frac{1}{\Phi} \left[\theta_X^{\sigma_y} + (1 - \theta_X)^{\sigma_y} (1 + t_E)^{1-\sigma_y} \right]^{1/1-\sigma_y} \tag{24}$$

Demand for the intermediate goods X and E can be derived by solving the cost minimization problem

$$\min_{X,E} \quad X + P_E \cdot E \qquad s.t. \quad f(X, E) \geq Y \tag{25}$$

with resulting demand functions of the form

$$X = \frac{Y}{\Phi} \cdot \left(\theta_X p_Y \right)^{\sigma_y} \tag{26}$$

$$E = \frac{Y}{\Phi} \cdot \left(\frac{(1 - \theta_X) p_Y}{1 + t_E} \right)^{\sigma_y} \tag{27}$$

Finally, market clearance implies that

$$Y = \frac{\sum_h \omega_h L_h + I_h}{p_Y} + G \tag{28}$$

The key complexity in representing this model is writing down a system of equations which represents $\tilde{V}^h$, the welfare level which agent h believes she could receive if she were able to set taxes herself. That is, consumers compute their counterfactual level of utility by choosing the taxes on labor and the dirty good. We make the following assumption about consumers' counterfactual utility problem:

a. Consumers know their own utility function and budget constraint.

b. They recognize the public budget constraint and take into account the connection between tax revenue and the level of provision of public goods.

c. While consumers recognize their own and the public budget constraints, we assume consumers are limited to a partial equilibrium framework of how policy instruments

19

affect prices for goods and factors, aggregate levels of demand and supply, and thus their own welfare.

d. Their model is *partial equilibrium* in the following dimensions:

    i. General equilibrium feedbacks on aggregate GDP are ignored, yet the counterfactual takes into account the elasticity of demand for energy and the uncompensated elasticity of labor supply by various households.

    ii. Energy demand is assumed isoelastic and depends on the level of economic output and energy prices:

$$E = \phi_E Y P_E^{\sigma} \tag{29}$$

    iii. Likewise, compensated labor supply is assumed isoelastic:

$$L = \phi_L P_L^{\eta} \tag{30}$$

The value of η which determines counterfactual labor supply is assumed consistent with the labor supply response from each of the household types:

$$\eta = \frac{\sum_h LS_h * \mathcal{E}_h}{\sum_h LS_h} \tag{31}$$

where

$$\mathcal{E}_h = \sigma_u + (1 - \sigma_u) \left(1 - P_L^{\theta_h^L} \right) \tag{32}$$

Since this is a relatively small departure from full rationality, consumers' counterfactuals are almost, but not quite correct. Naturally, larger errors would lead to much larger deviations between actual and consumer-computed counterfactual utilities.

We solve the model in calibrated share form and compute welfare, prices and quantities relative to a benchmark defined by $t_L = 0.5$, $t_E = 1$ (i.e. a 50-percent tax on labor, and a 100-percent tax on the dirty good). We separate consumers into three distinct groups:

- The "Green Young" have a high labor endowment $\bar{L}_h$ but little exogenous income I_h and represent the workforce. They have a relatively low preference for the public

good, but a high preference for environmental quality.

- The "Green Old" represent a consumer group with high non-labor income I_h and a high preference for environmental quality and the public good, but which does not participate in the labor market (technically, their shadow value of time exceeds the wage rate due to a low labor endowment).

- The "Brown Old" are similar to the Green Old except that they do not value environmental quality as highly.

The top part of Table 1 shows the endowments and taste parameters for these three consumer groups.

## 3.2 Results

Table 1: Model parameters and results

	Brown Old	Green Old	Green Young	Total
Parameters				
L_h	3	3	16	16
I_h	8	8	0	0
V_G	3	3	0.25	
V_Q	0.25	2	2	
EV (% of income)				
A' $\rightarrow$ A	0.00	1.65	-5.52	-3.87
A' $\rightarrow$ B	-1.44	3.15	-1.99	-0.29
A' $\rightarrow$ Fixed point	1.28	4.07	32.29	
A' $\rightarrow$ Utility maximum	3.10	4.70	32.40	
Welfare weights	0.47	0.47	0.06	1.00
Welfare change				
A' $\rightarrow$ A	0.00	0.77	-0.33	0.45
A' $\rightarrow$ B	-0.68	1.48	-0.12	0.68

Additional basic model parameters are $\sigma_u = \sigma_y = 0.5$, $\theta_X = 0.85$ and $\gamma = 2$. All other parameters are calibrated to the benchmark defined by $t_L = 0.5, t_E = 1$.

Figure 4 shows the numerical version of Figure 1. The corners of the Pareto-optimal policy space are again the true utility maxima for the three consumer groups, whereas the counterfactual utilities (as computed from A$'$) are marked by the triangles. In the figure, the coercion constraint is binding for the Brown Old consumer group, but not for the others.

The coercion-constrained equilibrium A$'$ is located at the tangency between the iso-coercion curve and the social indifference curve. This is the benchmark point from which

Figure 4: Coercion-constrained equilibrium
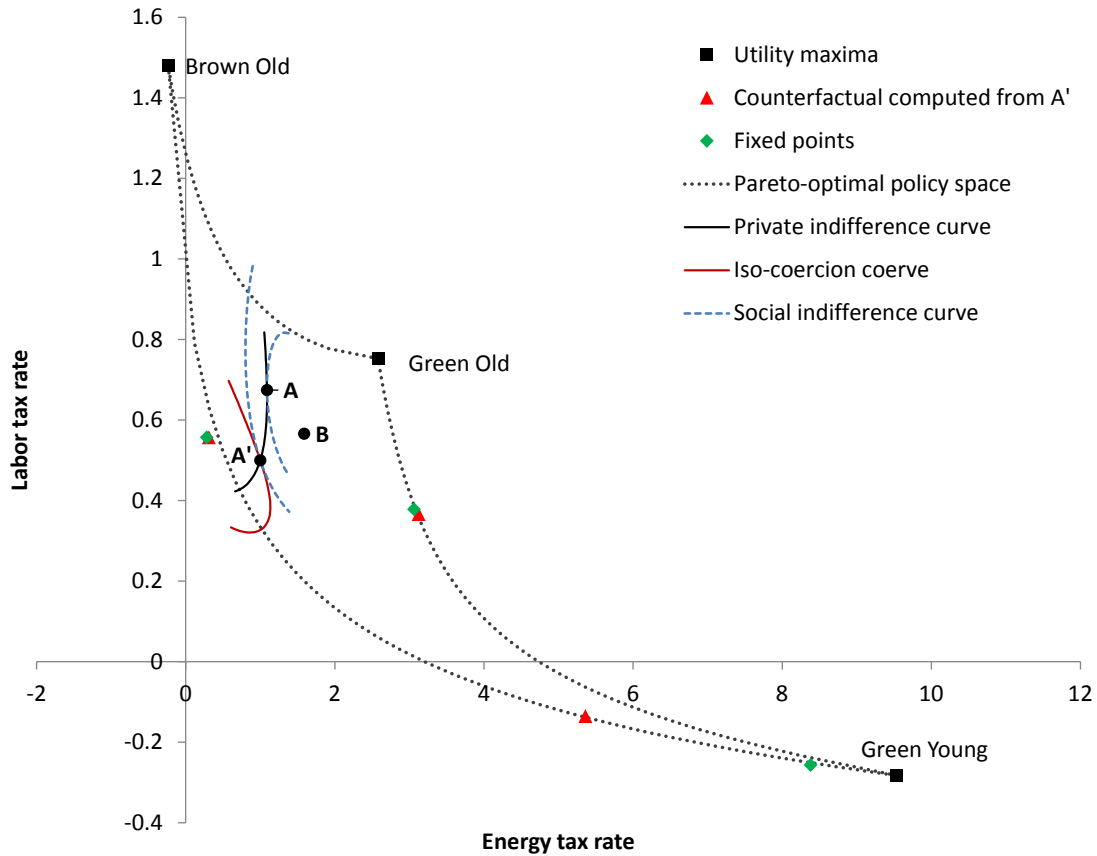
the non-basic model parameters are calibrated. Fixing counterfactual utility at the level computed at A′ would rotate the iso-coercion curve such that it coincides with the indifference curve going through point A′. The "error-corrected" equilibrium is located at point A, at the tangency between this indifference (and also iso-coercion) curve and a higher social indifference curve. If the coercion constraint were dropped altogether, the equilibrium would be located at point B, the unconstrained social welfare optimum.

The figure also shows the counterfactual fixed points as defined in (17). These are the points where consumers' counterfactual coincides with current policy, such that they would not want to move policy elsewhere. These fixed points do not coincide with consumers' true utility maxima due to the engogenous nature of the counterfactual: Although their utility would increase when moving the equilibrium towards their true optima, the errors are such that the perceived optimum, as computed from the fixed point, is given by the tax combination at the fixed point.

If the fixed points are located outside the policy space, an inefficient solution as discussed in the context of Figure 2 arises if the relevant consumer's coercion constraint is sufficiently tight: As the constraint approaches zero, the solution approaches the fixed

point and thus outside the Pareto-optimal policy space. Note that with a locally correct counterfactual, the fixed points are located at the true utility maxima, such that no inefficiency can arise.

Because of the scale-invariance of the utility function, we derive welfare statements based on willingness to pay measures. The middle panel of Table 1 show the equivalent variation, as a percentage of income, associated with moving from the benchmark point A′ to either point A, B, or consumers' fixed points and true utility maxima.[16] Multiplying these equivalent variations by the welfare weights that are implicit in the benchmark solution yields a measure for the change in social welfare associated with these policy changes (note that consumers all have the same income in our numerical model), which are shown in the bottom panel. The coercion constraint reduces social welfare by a (socially weighted) average of 0.68 percentage points of consumers' income. Of this welfare loss, 0.45 percentage points of income are attributable to the erroneous nature of the counterfactual, whereas only 0.24 percentage points are due to the unavoidable loss associated with imposing a coercion constraint in the first place. This implies that even a relatively mild error, such the one that we impose by consumers neglecting general-equilibrium effects when computing their counterfactual, can significantly increase the welfare loss associated with limiting economic coercion.

4   Conclusions

In this paper we posit that people compare their actual utility, which is conditional on policy decisions that are exogenous to them, to a hypothetical situation where they are able to define policy themselves. We call the utility that people would derive in this hypothetical world their "counterfactual utility". The larger the difference between actual and counterfactual utility, the more people feel coerced by the economic reality they observe.

Economic coercion is unavoidable if people differ with respect to their tastes or earning abilities, and it therefore cannot be a policy objective to eliminate economic coercion altogether. However, imposing a limit on the maximally acceptable level of economic coercion for particular consumer groups can be justified for several reasons, such as the

---

[16]The equivalent variation is the amount of income that has to be given to the consumer at the benchmark point A′ in order to make him indifferent between remaining there and moving to A or B. It is the monetized utility difference between A′ and A (or B), computed at the price level associated with A′. Note that the fixed points and utility maxima differ by consumer, whereas points A and B are common equilibria.

protection of minorities, the wish to preempt socially damaging action from individuals who feel disenfranchised, or (in a political economy-context) to secure the favor of swing-voters. While governments typically employ negotiations to settle labor disputes and law enforcement to prevent and punish illegal activities, mitigating disenfranchisement ex-ante by limiting economic coercion may be worthwhile.

The effect of limiting economic coercion depends crucially on whether people make mistakes when computing their counterfactual utilities. If they are completely rational and fully informed, their counterfactual utility becomes independent of actual policy decisions. Limiting coercion is then equivalent to increasing the welfare weight of particular consumers at the expense of others. The resulting equilibrium is different to the coercion-unconstrained case, but it is always allocationally efficient. The associated reduction in social welfare can be interpreted as the price society pays to reduce the risk for social strife below an acceptable level.

However, if consumers are not able to solve the full general equilibrium model, their counterfactual utility becomes endogenous to current policy. Considering the complexity of modern economies and the multitude of interactions between different policies, prices and levels of demand and supply, an assumption of limited rationality seems natural. We show that the welfare loss under incorrect counterfactuals exceeds the corresponding welfare loss if the counterfactuals were fixed. In other words, consumers' errors give rise to a welfare loss beyond the unavoidable loss associated with limiting coercion under full rationality. In our numerical model where we assume relatively consumer minor errors, the share of the welfare loss that is due to the error is about two-thirds. With a larger error, for example if consumers do not respect the government budget constraint or have a limited understanding about the available technology, this share would presumably increase.

If consumer errors are sufficiently large and the coercion constraint sufficiently tight, the coercion-constrained policy outcome may even be allocationally inefficient in the sense that it could not be reproduced by any set of nonnegative welfare weights (in the language used in this paper, the solution lies outside the Pareto-optimal policy space). We derive a necessary condition for this to occur.

Limiting economic coercion under limited rationality can be interpreted as a particular form of paternalism, to which we refer as *counterfactual* paternalism. By this we mean

that there exist policy choices that would increase the utility of the coercion-constrained consumer as well as social welfare. However, changing policy in this direction would lead the consumer to adjust his counterfactual in a way that would lead to a violation of the coercion constraint and thus reduce his perceived welfare. In other words, the social planner does not have superior knowledge about consumers' actual, but about their counterfactual utility, and he uses this knowledge to ensure that a pre-defined level of social discontent is not exceeded.

Naturally, our results depend on the particular error that we assume consumers to make. Although the error used in the numerical application seems plausible to us, it is admittedly ad-hoc. To generalize the results and develop more intuition about the importance of rationality on coercion, it would be worthwhile to examine various classes of errors and their effect on the coercion-constrained equilibrium, preferably in the context of a model that is calibrated to real economic data. This would allow for an assessment as to which types of common mis-perceptions the government should aim to reduce in order to limit economic coercion to some groups at the lowest social cost. Conversely, the results could be used to inform policy about the social cost of protecting special interest groups, conditional on the mis-perceptions believed to be present in those groups.

References

Baumol, William J (2004). "Welfare Economics and the Theory of the State." in *The Encyclopedia of Public Choice*. Springer: 937–940.

Bernheim, B. Douglas and Antonio Rangel (2007). "Behavioral Public Economics: Welfare and Policy Analysis with Nonstandard Decision-Makers." in Peter Diamond and Hannu Vartiainen eds. *Behavioral Economics and its Applications*. Princeton University Press: 7–77.

Boadway, R. (2014). "Discussion: The Role of Coercion in Public Economic Theory." in *Coercion and Social Welfare in Public Finance*. Cambridge University Press: 195–200.

Buchanan, James M. and Gordon Tullock (1962). *The Calculus of Consent. Logical Foundations of Constitutional Democracy*. University of Michigan Press.

Congleton, R. (2014). "Coercion, Taxation, and Voluntary Association." in *Coercion and Social Welfare in Public Finance*. Cambridge University Press: 91–116.

Gamson, William A. (1961). "A Theory of Coalition Formation." *American Sociological Review* 26(3): 373–382.

Ledyard, J. (2014). "Non-Coercion, Efficiency, and Incentive Compatibility in Public Goods." in *Coercion and Social Welfare in Public Finance*. Cambridge University Press: 143–159.

Lindahl, Eric (1919). "Just Taxation, a Positive Solution." in R. Musgrave and A. Peacock eds. *Classics in the Theory of Public Finance (1958)*. MacMillan: 168–176.

Martinez-Vazquez, Jorge and Stanley L Winer (2014). *Coercion and Social Welfare in Public Finance: Economic and Political Perspectives*. Cambridge University Press.

Skaperdas, S. (2014). "Proprietary Public Finance: On Its Emergence and Evolution Out of Anarchy." in *Coercion and Social Welfare in Public Finance*. Cambridge University Press: 60–81.

Skarpedas, Stergios (1992). "Cooperation, Conflict, and Power in the Absence of Property Rights." *American Economic Review* 82(4): 720–739.

Wallis, J. (2014). "The Constitution of Coercion: Wicksell, Violence, and the Ordering of Society." in *Coercion and Social Welfare in Public Finance*. Cambridge University Press: 29–59.

Wicksell, Knut (1896). "A New Principle of Just Taxation." in R. Musgrave and A. Peacock eds. *Classics in the Theory of Public Finance (1958)*. MacMillan: 72–118.

Winer, S., G. Tridimas and W. Hettich (2014). "Social Welfare and Coercion in Public Finance." in *Coercion and Social Welfare in Public Finance*. Cambridge University Press: 160–194.