Appendix (not for publication) for "Tax Reform and the Political Economy of Tax Base"

Ethan Ilzetzki*

This Draft June 27, 2017.

A.1 Proofs

Throughout, let

$$T(j) \equiv 1 - \tau(j); \qquad \hat{T} \equiv 1 - \hat{\tau}; \qquad \theta \equiv \frac{1-\tau}{1-\hat{\tau}} \leq 1.$$

The superscript T denotes a citizen who produces a taxed good and E denotes a citizen with a tax exemption.

A.1.1 Proposition 1

Using (7) in the main text we have

$$\frac{\partial u^E}{\partial k} = \frac{\partial u^E}{\partial \hat{T}} \frac{\partial \hat{T}}{\partial k},$$

for $k \in \{\tau, f\}$.

*Department of Economics, London School of Economics; Centre for Macroeconomics; and CEPR.

And using (9) in the main text, we have

$$\frac{d\log\rho}{d\tau} = \frac{1}{\tau} - \frac{\varepsilon-1}{T} + \frac{\eta-\varepsilon+1}{\hat{T}}\frac{\partial\hat{T}}{\partial\tau} \quad \text{and}$$
$$\frac{d\log\rho}{df} = \frac{1}{f} + \frac{\eta-\varepsilon+1}{\hat{T}}\frac{\partial\hat{T}}{\partial f}.$$

In addition:

$$\frac{\partial\hat{T}}{\partial\tau} = -f\theta^{\varepsilon-2}$$
$$\frac{\partial\hat{T}}{\partial f} = -\frac{1-T^{\varepsilon-1}}{(\varepsilon-1)\,\hat{T}^{\varepsilon-2}}f\theta^{\varepsilon-2}.$$

Using these equations, and the definition of the marginal cost of public funds, it can be shown that for exempt citizens,

$$MCPF^{\tau}(E) > MCPF^{f}(E) \tag{A.1}$$

is equivalent to

$$1-\tau-(\varepsilon-1)\,\tau-(1-\tau)^{\varepsilon} < 0,$$

if (10) from the main text holds. This latter inequality holds for all $\tau > 0$, so that exempt firms always prefer tax base increases to tax rate increases, as long as this does not change their tax status.

Turning to taxed citizens, it is easy to show that

$$\frac{\partial u^T}{\partial\tau}\Big/\frac{\partial u^T}{\partial f} > \frac{\partial u^E}{\partial\tau}\Big/\frac{\partial u^E}{\partial f},$$

i.e., citizens whose firms bear the brunt of taxation find increases in statutory taxes even more costly relative to base-broadening measures than do citizens with tax exemptions. It then follows directly that $MCPF^{\tau}(T) > MCPF^{f}(T)$, i.e. taxed citizens also prefer broadening the base to increasing statutory rates.

If $MCPF^{\tau}(j) > MCPF^{f}(j)$ for both the taxed and the exempt, it must be the case that the optimal tax base for every citizen is $f = 1$, keeping the tax status of the citizen in question unchanged.

A.1.2 Proposition 2

Define

$$U^j(f,g) \equiv \max_{\tau} \left\{ u^j(f,\tau) + \psi(\rho(f,\tau) - g) \right\}, \tag{A.2}$$

giving the utility of a citizen j if the tax base is f and the statutory rate is chosen to raise g units of revenues.

At a tax base of f and public good needs of g, the individual value of a tax exemption is equal to

$$\Delta^E U(f,g) \equiv U^E(f,g) - U^T(f,g):$$

the difference between the utility of the exempt and the taxed.

The value of a tax break and f Applying the envelope theorem to (A.2)

$$\frac{\partial U^j(f,g)}{\partial f} = \frac{\partial \rho(f,\tilde{\tau}(f,g))}{\partial f} \left(MCPF^{\tau}(j) - MCPF^{f}(j) \right),$$

where $\tilde{\tau}$ is the statutory tax rate that raises g units of revenues when the tax base is f. Then

$$\frac{\partial \Delta^E U(f,g)}{\partial f} < 0$$

iff

$$MCPF^{\tau}(T) - MCPF^{f}(T) > MCPF^{\tau}(E) - MCPF^{f}(E).$$

With some derivation, it can be shown that this is equivalent to

$$\frac{\partial \log u^T(f,\tau)}{\partial \hat{T}} \left[\frac{\partial \hat{T}}{\partial f} \frac{\partial \rho(f,\tau)/\partial \tau}{\partial \rho(f,\tau)/\partial f} - \frac{\partial \hat{T}}{\partial \tau} \right] - \frac{\partial \log u^T}{\partial \tau} >$$
$$\frac{\partial \log u^E(f,\tau)}{\partial \hat{T}} \left[\frac{\partial \hat{T}}{\partial f} \frac{\partial \rho(f,\tau)/\partial \tau}{\partial \rho(f,\tau)/\partial f} - \frac{\partial \hat{T}}{\partial \tau} \right].$$

This inequality holds because

$$\frac{\partial \log u^T(f,\tau)}{\partial \hat{T}} > \frac{\partial \log u^E(f,\tau)}{\partial \hat{T}}, \tag{A.3}$$

while

$$\frac{\partial \hat{T}}{\partial f} \frac{\partial \rho(f,\tau)/\partial \tau}{\partial \rho(f,\tau)/\partial f} - \frac{\partial \hat{T}}{\partial \tau} > 0,$$

and

$$\frac{\partial \log u^T}{\partial \tau} < 0. \tag{A.4}$$

Therefore the value of a tax break is decreasing in f.

The value of a tax break and g Applying the envelope theorem to (A.2)

$$\frac{\partial U^j(f,g)}{\partial g} = \psi = MCPF^\tau(j).$$

Then

$$\frac{\partial \Delta^E U(f,g)}{\partial g} > 0$$

iff

$$MCPF^\tau(T) > MCPF^\tau(E),$$

which is equivalent to

$$\frac{\partial \log u^T(f,\tau)}{\partial \hat{T}} \frac{\partial \hat{T}}{\partial \tau} + \frac{\partial \log u^T(f,\tau)/\partial \tau}{\partial \hat{T}/\partial \tau} > \frac{\partial \log u^E(f,\tau)}{\partial \hat{T}},$$

which holds because (A.3), (A.4) and $\frac{\partial \hat{T}}{\partial \tau} < 0$ hold.

The value of a tax break and ε Using (7) from the main text, the value of a tax exemption can alternatively be written as

$$\Delta^E U(f,g) = \left(\frac{z\hat{T}}{\mu}\right)^{\eta+1} (\mu - 1) \frac{1 - T^\varepsilon}{\hat{T}^{\varepsilon-1}}.$$

Then

$$\frac{\partial \log\left(\Delta^E U(f,g)\right)}{\partial \varepsilon} = -\log \hat{T} - \frac{T^\varepsilon \log T}{1 - T^\varepsilon} > 0,$$

because $\hat{T}$ and T are both smaller than 1.

The value of a tax break and η

$$\frac{\partial \log\left(\Delta^E U(f,g)\right)}{\partial \eta} = \log\left(\frac{\hat{T}z}{\mu}\right)$$

This is greater than zero if the equilibrium real after-tax real wage ($\hat{T}w$) is greater than one, i.e. if aggregate demand is increasing in the Frisch elasticity.

A.1.3 Proposition 3

Put formally, we are stating that for any g, there exists a value $f^R \in (0,1)$, such that for all $\tilde{f} \leq f^R$

$$U^T(1,g) \geq U^E(\tilde{f},g) \tag{A.5}$$

or

$$\rho(\tilde{f},\tilde{\tau}) < g$$

for all $\tilde{\tau}$, where $U(f,g)$ is given by (A.2). In words, there exists a cutoff tax base f^R, below which exempt citizens are better off with tax reform at $f = 1$ (at which they are taxed) than being at any feasible narrower tax base $\tilde{f} < f^R$.

The proof relies on the fact that the function $U^E(f,g)$ is increasing in f for all $f \in [0,1]$. In addition, at one extreme ($f = 1$), $U^E(1,g) > U^T(1,g)$, so that (A.5) is violated for sufficiently high values of f. This is true because a tax exemption gives positive utility, all else equal. At the other extreme, if f is sufficiently low, $U^E(f,g)$ goes to zero, while $U^T(1,g) > 0$, so that (A.5) must hold. With $U^E(f,g)$ increasing and spanning values above and below $U^T(1,g)$, there must be a value of $f = f^R$ at which these two utilities are equal. Moreover, citizens are better off with tax reform if and only if $f > f^R$.

I now demonstrate that it is the case that for sufficiently low values of f, $U^E(f,g)$ goes to zero. This is because the slope of $U^E(f,g)$ becomes infinite as f changes, for sufficiently low values of f. In the proof of proposition 2, I showed that

$$\frac{\partial U^E(f,g)}{\partial f} = \frac{\partial \rho(f,\tilde{\tau}(f,g))}{\partial f}\left(MCPF^{\tau}(j) - MCPF^{f}(j)\right) > 0.$$

For any feasible value of f, $\frac{\partial \rho}{\partial f}$ and $MCPF^f(j)$ are both strictly positive and finite. The numerator of $MCPF^{\tau}(j)$ is positive and finite, while its denominator goes to zero as f approaches its lowest feasible value. This last statement is true because the statutory tax rate must be set at its revenue-maximizing rate as f approaches its lowest feasible value. At the revenue maximizing rate, $\frac{\partial \rho(f,\tau)}{\partial \tau} = 0$, by definition. Thus $\frac{\partial U^E(f,g)}{\partial f}$ goes to infinity for sufficiently low values of f.

Below some threshold value of f it is of course the case that revenues of g are no longer feasible. It still remains the fact that no *feasible* policy exists with a narrower base that makes the exempt better off than under reform. (One might think of such cases as reform being induced due to feasibility rather than desirability to lobbyists.)

A.1.4 Proposition 4

It is easy to show that if f^R is determined by feasibility, i.e. f^R is the narrowest tax base at which g is feasible, then f^R is increasing in g. In this case, τ is set at its revenue-maximizing rate at f^R, given by

$$\frac{1-\bar{\tau}}{\bar{\tau}} = (\eta - \varepsilon) f\bar{\theta}^{\varepsilon} + \varepsilon, \quad \text{where} \qquad (A.6)$$
$$\bar{\theta} = \frac{1-\bar{\tau}}{f(1-\bar{\tau}) + 1 - f}.$$

It is easy to show that at the revenue maximizing rate $\bar{\tau}$, tax revenues are increasing in f, so that an increase in g requires an increase in f^R to remain feasible.

If, instead, f^R is determined by preferences, f^R is defined by

$$u^E\left(f^R, \tau^R\right) = u^T\left(1, \tau^1\right), \qquad (A.7)$$

where τ^R and τ^1 are the tax rates required to raise g in revenues, if the tax base is f^R and 1, respectively:

$$\rho\left(f^R, \tau^R\right) = \rho\left(1, \tau^1\right) = g \qquad (A.8)$$

This last equation simply maps τ^1 onto g with τ^1 naturally increasing in g. We can therefore conduct comparative statics with respect to τ^1, taken as an exogenous parameter with results for g following. $\{f^R, \tau^R\}$ are defined implicitly by (A.7) and the first equality in (A.8). Conducting comparative statics using these two equations, we obtain that

$$\frac{\partial f^R}{\partial \tau^1} = \frac{MCPF^E_\tau\left(f^R, \tau^R\right) - MCPF^T_\tau\left(1, \tau^1\right)}{MCPF^E_\tau\left(f^R, \tau^R\right) - MCPF^E_f\left(f^R, \tau^R\right)} \frac{\partial \rho\left(1, \tau^1\right)/\partial \tau}{\partial \rho\left(f^R, \tau^R\right)/\partial f}.$$

As long as we are on the correct side of the Laffer curve, the marginal revenue terms are positive. Also, proposition 1 states that the difference in marginal costs of public funds in the denominator is positive. Therefore $\frac{\partial f^R}{\partial g} > 0$ if and only if

$$MCPF^E_\tau\left(f^R, \tau^R\right) > MCPF^T_\tau\left(1, \tau^1\right). \qquad (A.9)$$

I now argue that this condition must hold for g sufficiently high. With g sufficiently high, we can make $\{f^R, \tau^R\}$ be arbitrarily close to the peak of the Laffer curve, but with τ^1 still away from the Laffer curve peak at $f = 1$. At the peak of the Laffer curve, marginal revenues are zero and $MCPF_\tau^E\left(f^R, \tau^R\right)$ goes to infinity. It must therefore be larger than the finite value of $MCPF_\tau^T\left(1, \tau^1\right)$. The inequality A.9 holds and f^R is increasing in g.

The second part of the proposition states that f^R is increasing in g for values of ε sufficiently high. Condition (A.9) can be rewritten as

$$\frac{1 + (\mu - 1)(\eta + 2 - \varepsilon) / \hat{T}\left(f^R, \tau^R\right)^{\varepsilon - 1}}{1 + (\mu - 1)(\eta + 2) T^1} >$$

$$\theta\left(f^R, \tau^R\right) \frac{1 - \frac{\tau^R}{1-\tau^R}\left[(\varepsilon - 1)\left(1 - f^R\theta\left(f^R, \tau^R\right)\right) + \eta f^R\theta\left(f^R, \tau^R\right)\right]}{1 - \frac{\tau^1}{1-\tau^1}\eta}.$$

Recalling that $\mu \equiv \frac{\varepsilon}{\varepsilon - 1}$, the left hand side of this equation can be made arbitrarily close to 1 for ε sufficiently high. The right hand side is always strictly less than one if $\varepsilon > \eta + 1$. This is because $\theta < 1$, $\tau^R > \tau^1$ and $f^R < 1$. Thus for ε sufficiently large, this inequality holds and f^R is increasing in g. Note that these are both sufficient–not a necessary–conditions. The relationship appears to hold for any values of ε, μ, and η.

A.1.5 Proposition 5

Follows directly from Proposition 3. If $f^q \leq f^R$, the policy $f = 1$ is Pareto improving compared to f^q. Hence zero compensation is required. In addition $f = 1$ is the social-welfare maximizing policy. Thus no other reform could improve welfare nor decrease the required compensation.

A.1.6 Proposition 6

Utility is

$$\tilde{u}^{PM}(f, \tau) = \left(\frac{\hat{T}z}{\mu}\right)^{\eta+1}\left[\frac{1}{1+\eta} + (\mu - 1)\left(\Lambda + (1 - \Lambda) L \frac{T(i)^\varepsilon}{\hat{T}^{\varepsilon - 1}}\right)\right]$$

For a given value of f, τ is determined by

$$U^{PM}(f, g) = \max_\tau \log \tilde{u}^{PM}(f, \tau)$$

s.t. $\log \rho(f,\tau) \geq \log g$. We are then asking whether there exists a value of g such that $U^{PM}(1,g) > U^{PM}(L,g)$. The proof then proceeds in a similar manner to that of Proposition 4, where the marginal cost of public funds using tax rates at $f = L$ must exceed that at $f = 1$, for g sufficiently high.

A.2 Empirical Evidence on the Timing of Corporate Tax Reforms

In this section I turn to the data and explore the timing of tax reforms in OECD countries. The main positive prediction from the previous section is that base-broadening tax reform occurs when the government faces high revenue needs. The normative analysis in section 4 recommends that reformers opt for large, rather than marginal, reforms. Whether a reform is large is in the eye of the beholder. Here, I define a reform as big if it frees up sufficient revenues to allow a reduction in statutory tax rates.

These two predictions are explored using data on corporate taxation collected by Kawano and Slemrod (KS, 2012). These data document legislative changes to the corporate tax base in 30 OECD countries from 1980 to 2004. The data include both high-income and emerging economies and episodes of rising and declining corporate tax rates. KS document changes in the tax base including changes in the generosity of investment credits, loss carry-forward rules, depreciation allowances, and others. I define a dummy variable *reform* that takes on a value of one if the tax base was broadened by any of the KS measures. In the reported regressions, I focus only on the tax base for domestic corporations as questions of international taxation go beyond the scope of this paper. Nevertheless, results are robust to including international tax legislation as well. Corporate taxation was not the main source of tax revenues for countries in the sample. However, data on the breadth of the income tax base or VAT base isn't readily available for a large sample of countries.

Corresponding to the theory, I measure public revenue needs by a single variable g: in this case government consumption as a percentage of GDP (source: World Bank).[1] In a dynamic context, however, governments can de-link current revenue needs from public good provision by borrowing. As I discuss below, using government revenues as a percentage of GDP (source: OECD), or debt to GDP (source: World Bank debt tables) leads to broadly similar results.

[1] While total government spending might be a more comprehensive measure of fiscal strain, the panel data coverage of this variable is much smaller.

Results from an OLS regression of the measure of fiscal strain on *reform* are shown in Table 1. The first column gives the result from an OLS regression of the *reform* dummy on government consumption as a fraction of GDP and shows that base-broadening reforms generally occur when spending is higher. This simple OLS regression does not have a causal interpretation, but suggests that base broadening reforms occur when public spending is high, consistent with the theory. Adding a control for GDP growth doesn't affect results and the coefficient on this control isn't statistically significant. Thus it is the state of public finances rather than the business cycle that drives reform.

The narrow focus on corporate taxation gives the government other (and larger) bases from which to draw revenues. It is therefore useful to have a measure that hones in on the fiscal strain to the corporate sector. Using the statutory tax rate as another measure of the fiscal needs of the government, particularly of the degree to which the government relies on corporate taxation for its fiscal needs, column 2 shows that governments do indeed broaden the corporate tax base when existing tax rates are high. This second regression also includes the *change* in the statutory corporate tax rate in the year of the reform. The negative coefficient suggests governments broaden the tax base when they lower statutory rates. This is consistent with the reform being large enough to free revenues to lower rates, despite high fiscal strain.

The fact that the tax base and statutory tax rates move in opposite directions is not an obvious accounting identity. If fiscal pressures are large, the government may be forced to increase both dimensions of tax policy simultaneously. In fact, this is what normative theories of the tax base, such as Yitzhaki (1979) would suggest. There, the tax base is narrow because of the administrative cost of enforcing a broader base. The optimal tax base equalizes the marginal cost of enforcing taxes with the marginal dead-weight losses due to the narrow tax base. In this normative literature, a government facing financing pressures would bear both these costs on the margin and increase the statutory rate while enforcing a broader tax base. The fact that governments cut statutory tax rates when broadening the tax base suggests that a different force is at play here. The result is consistent with the theory proposed in this paper, where the government needs the lower statutory tax rates to compensate losers from the base-broadening reform.

Conceptually, we are less interested in the cross-sectional aspect of the data, as our theory has more to say about when a given country passes reform than which country is more likely to pass reform measures. Indeed, country-specific factors may confound our results. The same political dysfunctionality that makes it hard

Table 1: Regression Results

	Dependent Variable = Reform					
	1	2	3	4	5	6
g/GDP	.009** (.004)	.010** (.004)	.009** (.004)	.032** (.014)	.008** (.003)	.007* (.004)
Corp. Tax Rate		.009*** (.002)	.007*** (.002)	.026*** (.009)	.008*** (.003)	.007* (.004)
Δ Tax Rate		-.017*** (.006)	-.018*** (.006)	-.060*** (.022)	-.020*** (.007)	-.014** (.006)
Country FE	NO	NO	YES	YES	YES	YES
Year FE	NO	NO	NO	NO	YES	YES
R^2	0.01	0.05	0.15	0.11	0.20	0.19
n	709	621	621	566	621	653

for a country to pass reform may make it harder for the government to raise revenues more generally. Similarly, a country unable to broaden the tax base may be forced to rely on higher statutory rates. The remaining columns therefore include country fixed effects. Column 3 shows that the results survive the inclusion of country fixed effects. The next two columns show that the results are robust to a probit specification (column 4), and inclusion of year (column 5) fixed effects.

Finally, in column 6, I replace government consumption with government revenues as a percentage of GDP (with country and year fixed effects). The results are robust to measuring fiscal strain in this way. The results are similar when using debt to GDP to measure fiscal strain, although in this case the result is statistically significant with year, but not country, fixed effects. This suggests that public indebtedness is a better predictor of which country will enact reform than when a specific country broadens its tax base. The result is nevertheless broadly consistent with the theory.

Certainly, one should be wary of simple correlations of this sort. These correlations are nevertheless suggestive that the theory presented in this paper is broadly consistent with a set of carefully-documented reforms.

A.3 Endogenous Government Spending

Consider a variant of the model where public good spending is endogenously determined. Specifically, let us augment individual preferences to include preferences over public spending:

$$u^j = x^j - \frac{\left(h^j\right)^{1+\frac{1}{\eta}}}{1+\eta} + \alpha^j v\left(g\right),$$

where $v\left(.\right)$ is a concave function and α^j reflects citizen j's preferences for public goods.

First note that preferences over the tax base and statutory rates are the same, regardless of what governs the amount of public spending to be financed. The first order condition

$$MCPF^{\tau}\left(j\right) \geq MCPF^{f}\left(j\right)$$

still holds and the proof of proposition 2 remains the same. Households always prefer to raise revenues with increases in the tax base rather than increases in the statutory rate, keeping their own tax status unchanged.

Second, for any desired public good needs g, it continues to be true that there is a critical tax base $f^R\left(g\right)$ below which an individual is willing to forgo her tax break in favor of tax reform at $f = 1$.

The first order condition governing desired public spending states

$$MCPF^{\tau}\left(j\right) = \alpha^j v'\left(g\right).$$

Thus now the value of f^R depends on and is increasing in α^j whenever f^R is increasing in g in the baseline model. We then have households with differing preferences over tax breaks vs. reform, with households with higher preferences for public goods more inclined towards tax reform (as a way to finance public good).

Turning to politics, once lobbying entry is determined, the analysis is almost identical to that in the benchmark model, except that the level of public good spending is determined by

$$MCPF^{\tau}\left(j\right) = \alpha^L v'\left(g\right),$$

where α^L is the average value of α^j among lobbyists. Tax reform occurs if $L > 1 - f^R\left(\alpha^L\right)$, i.e. if the average lobbyist prefers a tax base of $f = 1$ than her individual tax break.

In considering who chooses to lobby, special interests' choice is roughly the same as in the benchmark model. This is due to a free-riding problem with respect to the determination of public goods. Given the identity of the existing L lobbyists, an entering lobbyist has infinitesimal influence on public good spending. Special interests therefore enter based on the value they assign to a tax break (relative to the alternative of not having one *at the expected value of* g). This is the same across individuals. Thus lobbyists enter if the expected value of a tax break is equal to the fixed cost to lobbying and the lobby size L is determined as in the benchmark model.

In reality, special interests lobby for public goods, transfers, and regulatory policy. A large literature studies whether and how they overcome collective action problems in special interest organization. It is an interesting avenue for future research to consider how lobbying for tax breaks interacts with or competes with lobbying for other policies with budgetary implications. The simple framework studying here is, however, robust to the inclusion of endogenous public goods.

A.4 Partial Deductions

In the body of the paper, I assume that tax deductible goods are fully exempt. In reality, many deductions are partial. For example, mortgage interest is deductible only on the first $750,000 of a mortgage (down from $1 million prior to the tax reform passed in December 2017). In this appendix, I analyze the case of partial deductions and show that the basic mechanism holds with partial deductions as well. The following section explores the political economy implications of deductions being only partial.

Let χ denote the share of spending on goods $i > f$ that is tax deductible. The baseline case has a full dedcution, reflecting $\chi = 1$. The budget constraint of household j then becomes

$$\int_0^1 p\left(i\right) x^j\left(i\right) di \leq \left(1-\tau\right)\left(wh^j + \pi^j\right) + \chi\tau \int_f^1 p\left(i\right) x^j\left(i\right) di,$$

or

$$\int_0^1 p^c\left(i\right) x^j\left(i\right) di \leq wh^j + \pi^j,$$

with $p^c(i) \equiv \frac{p(i)}{T(i)}$ and where now $T(i) = T_E$ for $i \geq f$ and $T(i) = T$ for $i < f$. $T_E \equiv \frac{1-\tau}{1-\chi\tau}$ gives the effective net-of-tax rate for tax deductible goods, where $T^E = 1$ in the baseline case. The effective net-of-tax rate is now

$$\hat{T} \equiv \frac{1}{p^c} = \left(f T^{\varepsilon-1} + (1-f) T_E^{\varepsilon-1} \right)^{\frac{1}{\varepsilon-1}}.$$

Using this effective net-of-tax rate, all further analysis is identical to the baseline model, with the exception of revenues, which are now given as

$$\rho = \tau z h \frac{f T^{\varepsilon-1} + (1-\chi)(1-f) T_E^{\varepsilon-1}}{\hat{T}^{\varepsilon-1}}.$$

Figure A.1 now repeats Figure 2 from the main body of the paper for four values of the deductible share χ (1, 0.75, 0.5, and 0.25) in four panels. Each panel gives the utility of exempt citizens (in dashed lines) and taxed citizens (in solid lines) as a function of the tax base f. (I use the term "tax base" loosely, as the tax base now has two dimensions: the share of deductible goods $1-f$ and their deductible share χ.) The figures also highlight the value of f^R, the critical tax base below which exempt citizens would be better of if the forwent their exemptions collectively. The effects of χ on f^R are ambiguous, because of two competing forces. On one hand, partial deductibility makes a narrow tax base less costly, making tax reform less likely. (Partial deductibility flattens the utility functions in the figure, decreasing f^R). On the other hand, partial deductibility makes tax exemptions less valuable, making tax reform more likely. (Partial deductibility closes the gaps between the utility of exempt and taxed firms in the figure, increasing f^R). For the baseline parametrization used for figures in the paper ($\eta = 0.5$, $\varepsilon = 2$) these two forces roughly cancel out, so that regardless of the deductible share χ, the critical tax base for reform f^R is roughly the same. However, note that the term "tax base" is used loosely and reflects the share of household expenditure that obtains *any* tax deduction, rather than the share of household expenditure that is sheltered from taxation.

Figure A.2 summarizes this result by plotting the critical tax base f^R as a function of the magnitude of the exemption χ that tax deductible goods obtain. For the baseline parameter values, this is roughly flat.

In policy discussions of tax expenditures, the cost of these expenditures, or the breadth of the tax base, is typically measured in dollars or as a percent of tax revenues. In contrast, in this paper, we measure the tax base as the share of goods/household expenditure that is tax deductible. The baseline model, deductible goods are fully deductible ($\chi = 1$), so that the share of revenues lost is roughly $1 - f$. However, in the real world, taxes deductions are partial. Figure A.3 helps translate between the tax base in the theory with empirical measures of the tax base. For a tax base of $f = 0.7$, it shows in a solid line the share of revenues lost due to tax exemptions as a function of the deductible share χ. The dashed line gives a naive calculation of the share of revenues lost given by the product of the share of deductible goods and the deductible share $(1 - f)\,\chi$. The figure shows that the naive calculation is a reasonable approximation. (The overlap is imperfect because of increasing deadweight losses as the deductible share increases.)

A.5 Endogenous Tax System

In the baseline model, I have taken the tax system as given. This was partially for analytical convenience, as it allowed a simple dichotomy between statutory rates and the tax base. In practice there may also be administrative and simplicity rationales to have discrete exemptions and brackets. In this appendix I generalize the model to allow the government to choose both the number of exemptions provided and the magnitude of these exemptions.

Consider then a scenario where the government can choose the statutory tax rate τ, the tax base f, and the size of the exemption χ. The labor income tax applies to all workers as in the baseline model, but any expenditures on the measure $1 - f$ goods that are tax exempt entitle the consumer to a tax exemption at the rate of $\chi\tau$. I use the term "tax base" loosely to refer to f, but this is not precise. There are now two dimensions to the tax base: the number of exemptions and their size.

In the lobbying model, L citizens bear the fixed cost and lobby. They choose a policy that maximizes their joint utility (and does not discriminate between them). The policy maker sets $f = 1 - L$ and chooses τ and χ. This nests as a special case the possibility of $\chi = 0$, which is equivalent to a full tax reform $f = 1$.

Figure A.4 shows the policy that maximizes the utility of recipients of tax breaks (lobbyists) as a function of f $(= 1 - L)$. These are equi-revenue curves, with each curve drawn along a single value of g. Curves further to the right reflect higher revenues. The height of the curves on the y axis gives the magnitude of the exemption χ: the fraction of the statutory tax rate that is refunded to consumers. Obviously, the statutory tax rate also changes along these paths, but it isn't show in this figure.

As can be seen from the figure, $\chi = 1$ above a critical value of f, for every level of public good spending. Thus, as in the baseline model, there is a critical "tax base" above which 100% exemptions are provided in equilibrium, just as in the baseline model. Put in political terms, there is a critical number of lobbyists, above which lobbyists will themselves choose not to fully exempt their products from taxation. The smoothness afforded by the choice of exemption size χ means that the critical tax base at which lobbyists begin to show restraint is higher than in the baseline model. This can be seen by comparing the inflection points in figure A.4 with those of figure 3 in the main body of the paper. (The curves in both figures reflect the same values of g.) The values of f^R in figure 3 of the paper are smaller than the inflection points in figure A.4. With a continuous choice of the magnitude of the tax exemption, lobbyists need not wait (moving from right to left in these figures) until deadweight losses accumulate to make tax exemptions entirely undesirable. Instead, they can begin to reduce the size of exemption continuously.

Like f^R in the baseline model, the inflection point reflects a critical tax base below which "reform" of $\chi < 1$ occurs and above which lobbyists exempt themselves fully. As before, this critical tax base moves to the right as g increases across curves in the figure. Thus, as before, "tax reform" is more probable for higher public good needs.

As noted earlier, the notion of the tax base is multidimensional and therefore more diffuse in this more general setting. It is therefore not immediately apparent that policies to the left of the inflection points in the figure deserve the label "reform". While lobbyists show restraint in the magnitude of tax breaks, which is positive for welfare, this occurs at narrower tax bases, which are harmful for welfare. Figure A.5 parallels figure 3 in the main text. It shows the utility of taxed (solid lines) and exempt (dashed lines) citizens as a function of f, taking into account the equilibrium magnitude of the tax exemption shown in figure A.4. Vertical lines indicate the inflection point, to the right of which $\chi = 1$. Again, these are shown along equi-revenue paths. Looking first at the utility of exempt citizens, note that there is no point at which the exempt would forgo exemptions entirely. (All points on the dashed lines are above all points on the corresponding solid lines). The flexibility of continuous exemptions allows lobbyists to adjust the magnitude of exemptions gradually, so that "full reform" is only a limiting phenomenon. However, turning to the utility of taxed citizens, we see that the partial restraint shown by lobbyists deserves the title "reform". Above the critical mass, as the tax base narrows, lobbyists' restraint in the magnitude of tax exemptions they demand so that the utility of taxed citizens increases, despite the increased number of exemptions. This can be seen by the fact that the utility of taxed citizens is *declining* in the tax base below the critical value.

It remains true that "big bang" reform is less costly than gradual reform, when considering a full reform of the tax system–the full elimination of all tax exemptions. Starting from an arbitrary tax base f, compensating special interests one by one for the gradual removal of their exemptions would entail a cost equal to the integral of the difference between the dotted and the solid lines from f to 1. Compensating special interests for a big bang reform would entail a cost equal to the difference between the utility of the exempt at the tax base f and the utility of the taxed at $f = 1$. In contrast to the baseline model, big bang reform always necessitates some compensation. But this compensation is always smaller than the cost of picking off special interests one at a time. In fact, in this more general setting, the cost of compensating special interests may increase (initially) on the gradual reform path. Gradual reform may increase the resistance to reform making it more difficult (and perhaps less credible) to follow through on the reform path.

Turning to politics, figure A.6 shows the benefit of lobbying as a function of the measure of lobbyists L, once a specific value of g is realized. As in figure 5 in the main text, this is shown for two values of g. This figure differs from that in the main text only in that the benefit of lobbying declines continuously–rather than abruptly–as one crosses the critical value. As long as the exempt allocate full exemptions to themselves, the benefit of lobbying is identical to that in the main text. It is increasing in the number of lobbyists due to the strategic complementarities of tax exemptions. Once a critical threshold is passed, the lobby reduces the size of exemption χ and therefore the benefit of being tax exempt declines. This can be seen in figure A.5 by the narrowing gap between the utility of the exempt and the taxed, moving right to left, to the left of the vertical lines. As in the main text, higher values of g have two implications. First, they increase the value of a full exemption when a full exemption is allocated. And second, they decrease the maximum sustainable number of lobbyists before reform occurs.

Obviously, integrating over values of g would give expected benefits of lobbying of the same contour as those shown in figure 6 in the main text. The nature of political equilibrium is therefore identical to the one in the main text.

To summarize, the nature of the model changes little when allowing partial exemptions, which allows an endogenous determination of the tax system. Tax reform is more likely when public good needs are high and will involve a broadening of the tax base, in this case taking the form of a reduction in the generosity of tax exemptions, and a decrease in statutory rates. (This was not shown in the figures in this appendix, but is the resultant policy as the generosity of exemptions decreases.) Big bang reform is a better reform strategy than gradual reform.

A.6 Heterogeneous Productivities

In the analysis thus far, all households and firms were identical ex-ante. Any difference in outcomes was solely due to horizontal discrimination created through the political process. This section explores how the introduction of differences in firms' productivities affect the model's predictions.

Consider a model that differs from the one introduced in the body of the paper only in that each firm has an individual level of productivity $z(i) > 0$. Citizens remain identical in their preferences and productivity as workers, but now differ in the productivity of the firm they own. Let

$$z \equiv \left[\int_{i=0}^{1} z(i)^{\varepsilon-1}\, di\right]^{\frac{1}{\varepsilon-1}}, \tag{A.10}$$

denote average productivity.

Firms' profit maximization problem is roughly as before. Firms maximize profits, subject to the market wage w, but their production frontier is now

$$x(i) \leq z(i)\, h(i).$$

Firms set prices at a constant markup over marginal cost. If the price of the average good is normalized to 1, good i is priced at

$$p(i) = \frac{z}{z(i)}.$$

Define

$$Z(f) \equiv \int_{i=0}^{f} \left(\frac{z(i)}{z}\right)^{\varepsilon-1} di.$$

$Z(f)$ is a measure of the tax base when the measure of taxable goods is f. $Z(f)$ is strictly increasing in f, with a slope that is increasing in the productivity of the marginal firm introduced into the tax base:

$$\frac{\partial}{\partial f} Z(f) = \left(\frac{z(f)}{z}\right)^{\varepsilon-1}.$$

Using the definition of z in (A.10), $Z(1) = 1$ and $Z(f) \in [0,1]\ \forall f$.

The CPI is still given by

$$p^c = \frac{1}{1-\hat{\tau}},$$

if the effective tax rate is now defined as

$$1-\hat{\tau}\equiv\left(\left(1-\tau\right)^{\varepsilon-1}Z\left(f\right)+1-Z\left(f\right)\right)^{\frac{1}{\varepsilon-1}}.$$

If $z\left(i\right)=z\ \forall i$, as was the analysis in the benchmark model, then $Z\left(f\right)=f$ and the effective tax rate $\hat{\tau}$ is as in the homogeneous-productivity case. In the more general specification introduced here, the effective tax rate is a function of the tax base $Z\left(f\right)$, which is no longer necessarily equal to the measure of goods in the tax base f.

Pursuing the remainder of the analysis as before, the indirect utility of citizen j is now slightly modified as:

$$u^{j}=\eta^{\eta}\left(\frac{z\left(1-\hat{\tau}\right)}{\mu}\right)^{\eta+1}\left(\frac{1}{1+\eta}+\left(\mu-1\right)\left(\frac{z\left(j\right)}{z}\right)^{\varepsilon-1}\frac{\left(1-\tau\left(j\right)\right)^{\varepsilon}}{\left(1-\hat{\tau}\right)^{\varepsilon-1}}\right) \qquad \text{(A.11)}$$

Indirect utility differs from that of the benchmark analysis in two ways. First, the effective tax rate $\hat{\tau}$ is a function of the tax base $Z\left(f\right)$,which is not necessarily equal to the measure of taxable goods f. Second, the second term in the indirect utility function, giving the utility of the "entrepreneur" side of the household, is multiplied by $\left(\frac{z(i)}{z}\right)^{\varepsilon-1}$. Higher productivity translates into higher profits.

A glance at this more general indirect utility function highlights that much of the factors affecting citizens' preferences remain the same, even with heterogeneity in productivity. The "worker" side of the household is interested only in setting tax policy so as to minimize the effective tax rate $\hat{\tau}$. The entrepreneur side of the household again has conflicting interests: she too is concerned about the effective tax rate, due to its effects on aggregate demand, but gains a discrete benefit from securing a tax exemption. The new $z\left(j\right)$ term reflects that owners of more productive firms derive more income from profits and therefore have more to gain from securing tax breaks than do less productive firms.

The logarithm of revenues is now given by

$$\log\left(\rho\left(\tau,f\right)\right)=\log\tau+\log Z\left(f\right)+\eta\log\left(1-\hat{\tau}\right)+\left(\varepsilon-1\right)\log\left(\frac{1-\tau}{1-\hat{\tau}}\right)+\zeta\left(z,\eta,\varepsilon\right). \qquad \text{(A.12)}$$

This differs from the homogeneous productivity benchmark only in that the tax base $Z\left(f\right)$ is no longer equal to the measure of taxable goods f.

Citizens producing taxed goods prefer lower tax rates and a narrower tax base, as in lemma 1. The condition in that proposition, necessary and sufficient for sheltered citizens to be tax averse, is now complicated by the heterogeneous productivity term $z(i)$ and now reads

$$(1-\hat{\tau})^{\varepsilon-1} > \left(\frac{z(j)}{z}\right)^{\varepsilon} (\mu-1)(\varepsilon-\eta-2). \tag{A.13}$$

This is now a citizen-specific condition, delineating the part of the state space, in which citizen j is tax averse. If the assumption of tax aversion from the benchmark model holds, all citizens with below-average productivity prefer lower tax rates and a lower tax base in this case, but citizens with above-average productivity may prefer higher tax rates and a narrower tax base. To abstract from this case, let us modify the assumption so that *all* citizens are tax averse. As in the homogeneous productivity case, $\eta+1 > \varepsilon-1$ is sufficient for (A.13) to hold for all j. In contrast, if $\varepsilon-1 > \eta+1$ and the distribution of $z(j)$ is unbounded, there is necessarily some citizen that likes higher taxes. Assumption (A.13) is stronger than that in the benchmark model and implies a restriction on the upper tail of the productivity distribution.

With this assumption, we are still able to find a value $f^R(z(j))$ such that citizen j prefers tax reform at $f=1$ to her own tax exemption if and only if $f < f^R(z(j))$. In addition to the comparative statics of proposition 2 in the paper, it can be shown that willingness to pay for a tax break is increasing in $z(j)$, so that more productive firms are more willing (and generally more able) to pay for a tax break. $f^R(z(j))$ is increasing in $z(j)$.

Turning to politics, nothing substantive changes. It is interesting to note, however, that higher productivity firms are more willing to lobby for a tax break, all else equal, and thus are more likely to be represented in setting tax policy.

Of course, firms may differ in their political power in ways that are correlated or uncorrelated with their productivity. If we think of low costs to enter lobbying as a reflection of higher political power, it then become more difficult to pin down whether lobbyists are the most productive firms. The correlation between economic power (productivity, wealth) and political power (barriers to entry into lobbying) appears to be an important determinant the beneficiaries of tax exemptions. I leave this question to future research.

A.7 Median Voter Model

In this appendix, I explore prediction of the model in a slightly different setting. While lobbying is central to the politics of tax exemptions, voters are also influential in setting tax policy. This appendix considers the role of voters in the canonical Downsian median voter model. The policy space in this setting is infinitely dimensioned, a tax base f can provide a measure of $1 - f$ tax exemptions to an infinite set of permutations of citizens. To sharpen analysis, in this section we'll assume that citizens are ranked in order of "taxibility", so that higher a lower indexed citizen obtains a tax break only if all citizens with higher indexes also obtain tax exemptions. This ranking could be due to administrative ease in giving tax exemptions to certain goods rather than others or due to a ranking of political power, among other factors. This reduces the problem into two dimensions. Voters vote on a tax base f which provides tax exemptions to all citizens with indexes $j \geq f$. The statutory tax rate is then determined residually through the budget constraint.

Condorcet Winner Let's begin by searching for a Condorcet winner, i.e. a policy that would receive a majority of votes in a bilateral referendum against any other policy. A Condorcet winner exists if $f^R > \frac{1}{2}$, but not if $f^R \leq \frac{1}{2}$, as outlined in the following proposition.

Proposition 1 *If $f^R > \frac{1}{2}$, there exists a Condorcet winning policy at $f = 1$. If $f^R \leq \frac{1}{2}$, no Condorcet winner exists.*

Proof. When $f^R > 0.5$, all voters $j < f^R$ prefer $f = 1$ to any other policy, and this is the Condorcet winner. If $j^R < 0.5$, no Condorcet winner exists. Any policy proposal $f < 1$ is dominated by a slightly broader base: all citizens vote for such a measure except the small number of citizens whose tax exemption was eliminated. But $f = 1$ is dominated by any $f \in \left[f^R, \frac{1}{2}\right)$ as this is supported by a measure $1 - f > \frac{1}{2}$ of citizens, who prefer a tax exemption to $f = 1$, by the definition of f^R. ■

The intuition for the first part of the proposition is straightforward: If $f^R > \frac{1}{2}$, the median voter is part of the cohesive coalition for tax reform and this policy is implemented. If $f^R \leq \frac{1}{2}$, only a minority of voters have tax reform as their ideal policy. In a median voter model, there is no direct way to resolve the collective action problem among citizens $j \geq f^R$ to form a unique winning coalition. It is therefore possible construct a policy that a winning coalition of voters would support in favor of any other policy in a bilateral vote. All citizens prefer a broader base, as long as their own tax status is unaffected by this change. For any $f < 1$, it is possible to broaden the base in such a way that a majority of voters is unaffected; this majority would prefer this broader base. However, there is also a coalition that would prefer any $f \in \left[f^R, \frac{1}{2}\right]$ to $f = 1$, by the very definition of f^R. With $f^R < \frac{1}{2}$ this coalition is a majority.

The absence of a Condorcet winner in the case $j^R < \frac{1}{2}$ poses problems of equilibrium existence in pure strategies. We now turn to equilibrium in mixed strategies under a winner-take-all electoral system, where candidates maximize their probability of obtaining a majority of votes. The result is that $f = 1$ is the unique equilibrium if $f^R > \frac{1}{2}$, but tax reform occurs with a lower probability if $f^R < \frac{1}{2}$.[2]

Political Model There are two political candidates A and B that are not citizens of the economy described so far. Their sole objective is to maximize their probability of election.

The political game consists of two stages. In the first stage–the voting stage–the two candidates observe the revenue requirement g and formulate their strategy. A strategy for each candidate is a choice of a probability distribution $\phi^A\left(f^A\right)$ or $\phi^B\left(f^B\right)$, respectively, from which they draw their political platforms. Let $\Phi^A\left(f^A\right)$ and $\Phi^B\left(f^B\right)$ denote the corresponding cumulative distribution functions. A well-defined probability distribution has $\phi(f) \geq 0$ for all $f \in [0,1]$ and satisfies $\Phi(f) = 0$ and $\Phi(1) = 1$. Candidates simultaneously draw platforms from their distributions. The platform consists of a tax base f^A or f^B. The corresponding statutory tax rate is uniquely determined by the budget constraint. The candidates are fully committed to implement their platform if elected.

Voters observe the platforms f^A and f^B and vote sincerely for their preferred candidate. Each citizen has one vote. Indifferent voters randomize between the two candidates with equal probability. The candidate who receives a majority implements her proposed policy. If both candidates obtain the same vote share, each candidate's policy is implemented with probability $\frac{1}{2}$.

In the second stage–the economic stage–the economy proceeds as in the description of the economy in the paper. Citizens choose their labor supply and consumption, firms maximize profits, and citizens' payoffs are realized, given the tax policy set in the first stage.

Political Equilibrium Not surprisingly, when $f^R > \frac{1}{2}$, the Condorcet-winning policy of $f = 1$ is implemented with probability 1. When $f^R \leq \frac{1}{2}$, in contrast, a

[2] An analysis of proportional representation (PR), approximated as in Lizzeri and Persico (2001), with vote-share maximization, is available upon request. One important difference is that the probability of tax reform converges smoothly to 1 as j^R approaches $\frac{1}{2}$ from below under PR, but jumps discretely from zero to 1 under winner take all.

large set of equilibria exist, but they all have some common features.

In all equilibria there is a negligible probability that comprehensive tax reform $f = 1$ is implemented. There is a 50% probability that a narrow base in the $\left[f^R, \frac{1}{2}\right]$ range is implemented and a 50% chance that a broader base in the $\left[f^R + \frac{1}{2}, 1\right]$ range is implemented. The following proposition summarizes the characteristics of equilibrium. I focus on symmetrical equilibria where both candidates draw platforms from the same distribution $\phi(f) = \phi^A\left(f^A\right) = \phi^B\left(f^B\right)$ with CDF $\Phi(f)$.

Proposition 2 ***Political Equilibrium.*** *If $f^R > \frac{1}{2}$, both candidates propose $f^A = f^B = 1$ with probability 1 and each obtains a vote share of $\frac{1}{2}$. The unique equilibrium policy is $f = 1$.*
If $f^R \leq \frac{1}{2}$, a function $\phi(.)$ constitutes a symmetrical political equilibrium if and only if it has the following characteristics.
1) There is a 50% probability of proposing a tax base between f^R and $\frac{1}{2}$.
2) There is a 50% probability of proposing a tax base between $f^R + \frac{1}{2}$ and 1.
3) $\phi(.)$ has an identical distribution in $\left[f^R, \frac{1}{2}\right]$ as in $\left[f^R + \frac{1}{2}, 1\right]$: $\phi(f) = \phi\left(f + \frac{1}{2}\right)$ $\forall f \in \left[0, \frac{1}{2}\right]$
4) The function $\phi(.)$ does not contain any mass points.

Proof. The paragraphs that follow provide a proof that the conditions above are sufficient for an equilibrium. The proof of their necessity is available upon request. ■

The intuition of the first part of the proposition is simple. The policy $f = 1$ is the preferred policy of a measure f^R of the population. If $f^R > \frac{1}{2}$, proposing this policy gives a candidate a vote share of no less than 50% and thus a probability of no less than 50% of winning. This clearly dominates any other strategy.

Figure A.7 illustrates an equilibrium for the case $f^R \leq \frac{1}{2}$. The proposition states that $\phi(.)$ looks identical in the ranges $F^L \equiv \left[f^R, \frac{1}{2}\right]$ and $F^H \equiv \left[f^R + \frac{1}{2}, 1\right]$, with a cumulative 50% probability of drawing a policy in either of these ranges. There is a zero probability of drawing a tax base outside F^L or F^H.

To see why this constitutes an equilibrium, it is first useful to note that any proposal $f^A \geq f^R$ defeats another proposal $f^B \geq f^R$ if and only if f^A proposes a broader base, without removing tax exemptions from more than 50% of citizens, relative to f^B. Citizens whose tax status is the same under both proposals prefer a broader tax base, but citizens $j \geq f^R$ prefer a tax exemption to a broader base.

With a symmetrical distribution as in figure A.7, a platform f^A in F^H then defeats all proposals f^B drawn from F^H such that $f^B < f^A$. It also defeats all proposals f^B drawn from F^L that are *greater* than $f^A - \frac{1}{2}$. As the distributions F^H and F^L are symmetrical, this adds up to a 50% chance of a proposal f^A winning against a proposal f^B drawn from the distribution. A similar logic applies to policies f^A drawn from F^L.

No proposal outside F^L or F^H is a profitable deviation. A proposal $f^A < f^R$ loses against all proposals in F^L (and possibly some in F^H) and thus cannot give a vote share of more than 50%. Any proposal in $\left[\frac{1}{2}, \frac{1}{2} + f^R\right]$ loses to proposals in F^H and obtains 50% of the vote. The PDF $\phi(f)$ in figure A.7 is therefore an equilibrium.[3]

Characteristics of Political Equilibrium Predictions of the model are sharper when $f^R > \frac{1}{2}$ and the unique equilibrium is tax reform $f = 1$. As is often the case with mixed-strategy equilibria, predictions are slightly murkier when $f^R < \frac{1}{2}$. A comparison between the two does, nevertheless, deliver some insights.

When revenue needs cross the $f^R = \frac{1}{2}$ threshold, policy changes discontinuously from a narrow base–a comprehensive tax base of $f = 1$ is a zero probability event–to comprehensive reform. Below this threshold there is a 50% probability that a broad, but incomplete, tax base is proposed (in the F^H range) and as f^R approaches $\frac{1}{2}$, this proposal becomes closer to comprehensive tax reform (in the sense that it incorporates a base that approaches 1). Nevertheless, as long as $f^R \leq \frac{1}{2}$, there is always a 50% probability that a narrow base is chosen within F^L. This probability is eliminated discontinuously at $f^R = \frac{1}{2}$.

The probability of disagreement among political candidates also changes discontinuously at $f^R = \frac{1}{2}$. When $f^R \leq \frac{1}{2}$, the probability that both candidates propose the same platform is zero and there is a 50% chance that they propose platforms in different ranges (F^L vs. F^H). Neither political party is likely to propose comprehensive tax reform ($f = 1$), but it is likely that one political party will adopt the mantle of some tax reform measure (F^H), with the other defending tax exemptions (F^L). The latter cultivates special interests through tax exemptions, while the former calls for a broader base, hoping to alienate less than 50% of voters in the process. In contrast, when tax revenue needs increase to the tipping

[3] The nature of equilibrium echoes Lizzeri and Persico (2001), where a public good is provided with probability one if its value is large enough relative to the value of targeted transfers. Here, the "public good" is tax reform whose value relative to rents generated by tax exemptions is determined endogenously through the economic structure of the model.

point where f^R exceeds $\frac{1}{2}$, the nature of political changes. Tax reform becomes political consensus, and both political parties put forth comprehensive tax reform proposals. This occurs because the majority of voters have internalized the need for tax reform.

A.8 Legislative Bargaining

In this appendix, I explore results in a legislative bargaining framework, in the spirit of Baron and Ferejohn (1989). The measure-one population is divided into n districts, each represented by legislator. One of the n legislators is randomly selected to be the agenda setter; each legislator has equal probability of selection. All legislators observe g and the agenda setter proposes a feasible policy, which is enacted if $m-1$ other legislators favor it, with $m < n$ the size of the minimum winning coalition. If the proposal does not receive sufficient support, a different legislator is selected to be agenda setter and the game proceeds in the same fashion. If no proposal receives support in T legislative bargaining rounds, a default policy of $f = 1$ is enacted. Policies cannot discriminate among citizens within a district.

It is easy to see that equilibrium policy is $f = 1 - \frac{m}{n}$ if $f^R \leq 1 - \frac{m}{n}$ and $f = 1$ otherwise. The agenda setter will never propose tax exemptions to any more than $m-1$ other districts, as further tax breaks would narrow the tax base without any political benefit. Let us restrict attention to an equilibrium where the agenda setter's proposal passes in the first round. An equilibrium the utility-maximizing policy for the agenda setter, such that $m-1$ legislators are better off accepting the proposal than proceeding to the second legislative round. Consider first the case where $f^R \geq 1 - \frac{m}{n}$ holds. In this case, the coalition is better off allocating tax exemptions to all districts than to none (which would constitute tax reform). This is moreover the preferred policy of the agenda setter. Any district not receiving a tax exemptions will vote against any other proposal, as they have the hope of becoming agenda setter (with probability $\frac{1}{n}$) or coalition members (with probability $\frac{m-1}{n}$) in the second round, which gives them a tax exemption with positive probability. Thus obtaining m supporters for the proposal requires giving m district tax exemptions, leading to a tax base of $1 - \frac{m}{n}$. If, on the other hand, $f^R > 1 - \frac{m}{n}$, the coalition is better off setting $f = 1$. Knowing that tax reform would be passed in the second round if they refuse, they are no worse off accepting tax reform in the first round.

To conclude, a legislative bargaining model leads to a similar tipping point to reform. Here, the critical value is the size of the minimum winning coalition rather than the number of lobbyists as in the main text.

A.9 Generalized Political Preferences

This section analyzes policy for a reduced-form, but relatively general, political objective function, where social weights may be arbitrarily assigned across citizens. This "political planner" may be viewed as shorthand for a variety of political mechanisms that might lead to less than equal weights across the citizenry. As Persson and Tabellini (2002) show, lobbying power due to political organization, the additional political influence that swing voters may exert, and several other political mechanisms can be reduced to a social welfare function that puts higher weights on the more influential elements in society.

The analysis in this section demonstrates that two of the main positive predictions of the theory survive this more general setting. Higher government revenue needs lead to a more reformed tax system: one that is more efficient and horizontally equitable. Specifically, the equilibrium tax base is increasing in the government's revenue needs. For revenue needs that are sufficiently high, $f = 1$ is the political outcome. Thus higher revenue needs catalyze tax reform.

In addition, the statutory tax rate is *decreasing* in revenue needs. This is true for a large variety of weights in the planner's objective function. As in the benchmark model, the political planner (partially) compensates losers from the broader tax base through a reduction in tax rates.

It is also illuminating to consider a result from the benchmark model that does not survive this analysis. I restrict attention to a welfare function that is smooth across citizens' indexes. No citizen has a mass weight and there is no discontinuity in political power across (groups of) individuals. Due to the smooth nature of the welfare function, tax reform no longer takes a "big bang" form. Rather, tax policy evolves smoothly as tax revenues change. This highlights that the existence of pivotal actors who have discretely greater political influence than others (as in the median voter analysis in this document and in the benchmark model) is a necessary ingredient for the result that large tax reforms are more likely than marginal ones.

Consider an economy that is identical to the benchmark model in the body of the paper. Citizens' preferences are

$$u^j = \left(\frac{z\hat{T}}{\mu}\right)^{\eta+1}\left(\frac{1}{1+\eta} + (\mu - 1)\frac{T(j)^{\varepsilon}}{\hat{T}^{\varepsilon-1}}\right),$$

where $\hat{T} \equiv 1 - \hat{\tau}$ and $T(j) \equiv 1 - \tau(j)$. The political planner puts a finite weight of $\beta(j)$ on the utility of each citizen j, where it is assumed that $\beta(j)$ is monotonically increasing in j: citizens are ordered according to their political influence. Also,

$$\int_{j=0}^{1} \beta(j)\, dj = 1.$$

Let $\mathcal{B}(f)$ denote the cumulative weight that the planner puts on citizens $j \in [0, f]$, so that

$$\mathcal{B}(f) \equiv \int_{j=0}^{f} \beta(j)\, dj$$

and $\mathcal{B}(1) = 1$.

Given that $\mathcal{B}(j)$ is increasing in j, a planner choosing a tax base of f will always give tax exemptions to citizens $j \in [f, 1]$. The planner's policy choice can therefore be described as

$$\max_{\tau, f} \left(\frac{z\hat{T}}{\mu}\right)^{\eta+1} \left(\frac{1}{1+\eta} + (\mu - 1)\frac{\mathcal{B}(f)\, T^{\varepsilon} + 1 - \mathcal{B}(f)}{\hat{T}^{\varepsilon-1}}\right),$$

subject to the budget constraint $\rho(\tau, f) > g$. The resulting policy will satisfy

$$MCPF_{PP}^{\tau} \geq MCPF_{PP}^{f}, \tag{A.14}$$

with equality if $f < 1$, where PP represents "political planner" and the two terms are the marginal costs of public funds to this planner from rasing revenues with each of the tax instruments.

The numerical analysis that follows shows that the tax base broadens and the statutory tax rate decreases as g increases. Before turning to numerical results, I show analytically that these results holds at two polar values, as g approaches zero and as g approaches the maximal revenues that the tax system can bear (the peak of the Laffer curve).

Let's first look at the case $g \to 0$. It can be shown that

$$\lim_{\{f,\tau\} \to \{0,0\}} \left\{\frac{MCPF_{PP}^{\tau}}{MCPF_{PP}^{f}}\right\} = \lim_{\{f,\tau\} \to \{0,0\}} \left\{\tau \frac{\beta(f) T^{\varepsilon-1}}{1 - \beta(f) T^{\varepsilon}}\right\},$$

where the limit can be taken along any path. As $\tau \to 0$, $T \to 1$. As $f \to 0$, $\beta(f) \to \beta(0)$. This limit is therefore equal to zero unless $\beta(0) = 1$. The marginal cost of raising a unit of revenues via increases in the tax rate goes to zero faster than the marginal cost of raising the same unit of revenues via broadening the tax base. It must therefore be the case that the political planner wishes to set a non-zero tax rate even as the tax base approaches zero. At near-zero revenues, the political planner relies more on a high statutory tax rate and a narrow tax base to raise revenues. This result does not hold when $\beta(0) = 1$. but $\beta(0) < 1$ for any increasing political welfare function that does not weigh all citizens identically.

I now look at the case where $g \to \bar{g}$, where $\bar{g}$ is the maximal level of public spending that can be sustained for any $\{f, \tau\}$. The revenue-maximizing tax rate satisfies

$$\frac{1}{\tau} - \frac{\varepsilon - 1}{T} + \frac{\varepsilon - 1 - \eta}{T} f\theta^{\varepsilon - 1} = 0,$$

where $\theta \equiv \frac{T}{\hat{T}}$. The revenue-maximizing tax base satisfies

$$\frac{1}{f} + \left(\frac{\varepsilon - 1 - \eta}{\varepsilon - 1}\right)\frac{1 - T^{\varepsilon - 1}}{\hat{T}^{\varepsilon - 1}} \geq 0,$$

with equality if $f < 1$. This condition cannot hold with equality for any $\tau > 0$, so that $f = 1$ is the revenue maximizing tax base. As $g \to \bar{g}$, therefore $\frac{\partial \log(\rho)}{\partial \tau} \to 0$, while

$$\frac{\partial \log(\rho)}{\partial f} = \frac{1}{f} + \left(\frac{\varepsilon - \eta}{\varepsilon}\right)\frac{1 - T^{\varepsilon}}{\hat{T}^{\varepsilon}} > 0,$$

as revenue needs approach their maximal feasible value. Given that the marginal utilities $\frac{\partial u^{PP}}{\partial \tau}$ and $\frac{\partial u^{PP}}{\partial f}$ are strictly positive and bounded, it must be the case that for revenue needs that are sufficiently high

$$MCPF^{\tau}_{PP} > MCPF^{f}_{PP},$$

giving $f = 1$ as the optimal policy.

To summarize, we have the result that the tax base will be narrow for revenue needs that are sufficiently low and broad at $f = 1$ for revenue needs that are sufficiently high. This result is demonstrated in Figure A.8, which displays the equilibrium tax base (left axis) and statutory tax rate (right axis) as the revenue needs of the government increase (left to right along the X-axis). As described in the discussion above, when revenue needs approach zero, the political planner relies on a narrow base and a high tax rate.

Figure A.8 indicates that there is a monotonic shift to a broader base as revenue needs increase. Despite the greater revenue needs, the tax rate declines. This continues until the tax system is "reformed" at $f = 1$. At this point, the only feasible way to raise revenues is through increasing statutory tax rates.[4]

To summarize, the result that higher revenue needs will lead to a more efficient tax system together with lower statutory rates holds in this more general setting. Unlike the benchmark model there is no tipping point, at which policy jumps discontinuously to $f = 1$. This illustrates the importance of a group of citizens with discretely greater political power–a pivotal interest, an organized group, etc., in the benchmark model.

References

[1] Baron, David P. and John A. Ferejohn, 1989"Bargaining in Legislatures." *The American Political Science Review* 83:4.

[4]The results shown here are for $\varepsilon = 2$, $\eta = 0.5$ and welfare weights of $\beta(j) = 2j$. They were robust to any of a wide range of parameter values and welfare weight functions, with which I simulated the model.

Figure A.1: Utility of Taxed and Exempt Citizens and and Deductible Share

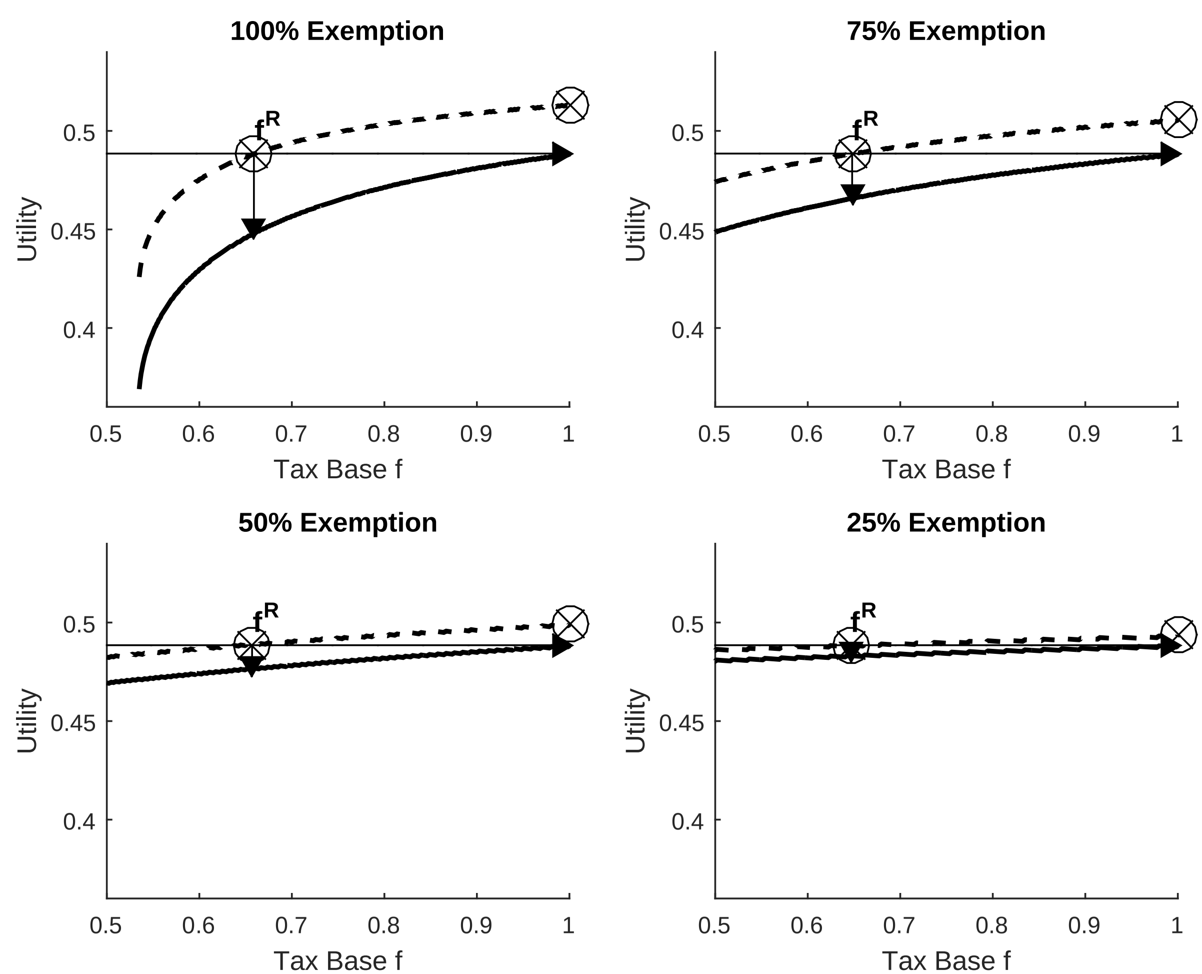

This figure shows the utility of owners of tax exempt firms (dashed lines) and taxed firms (solid lines) for four values of the share of expenditure on exempt goods that is tax deductible χ. The values are 1, 0.75, 0.5 and 0.25. Each panel also indicates the critical tax base f^R below which exempt firms would benefit from collectively forgoing their tax exemptions.

Figure A.2: Critical Tax Base for Reform and Deductible Share

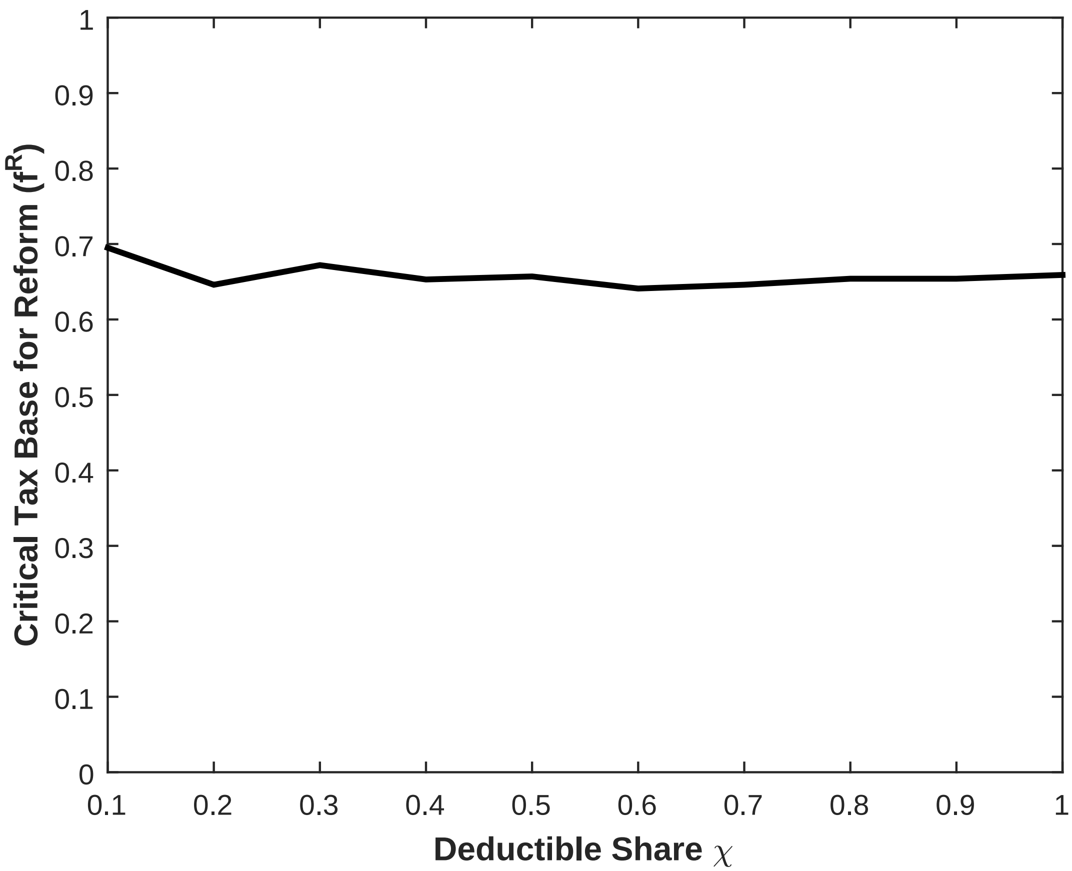

This figure shows the critical tax base f^R, below which exempt firms would benefit from collectively forgoing their tax exemptions, as a function of χ, the share of expenditure on exempt goods that is tax deductible.

Figure A.3: Revenues Lost and Deductible Share

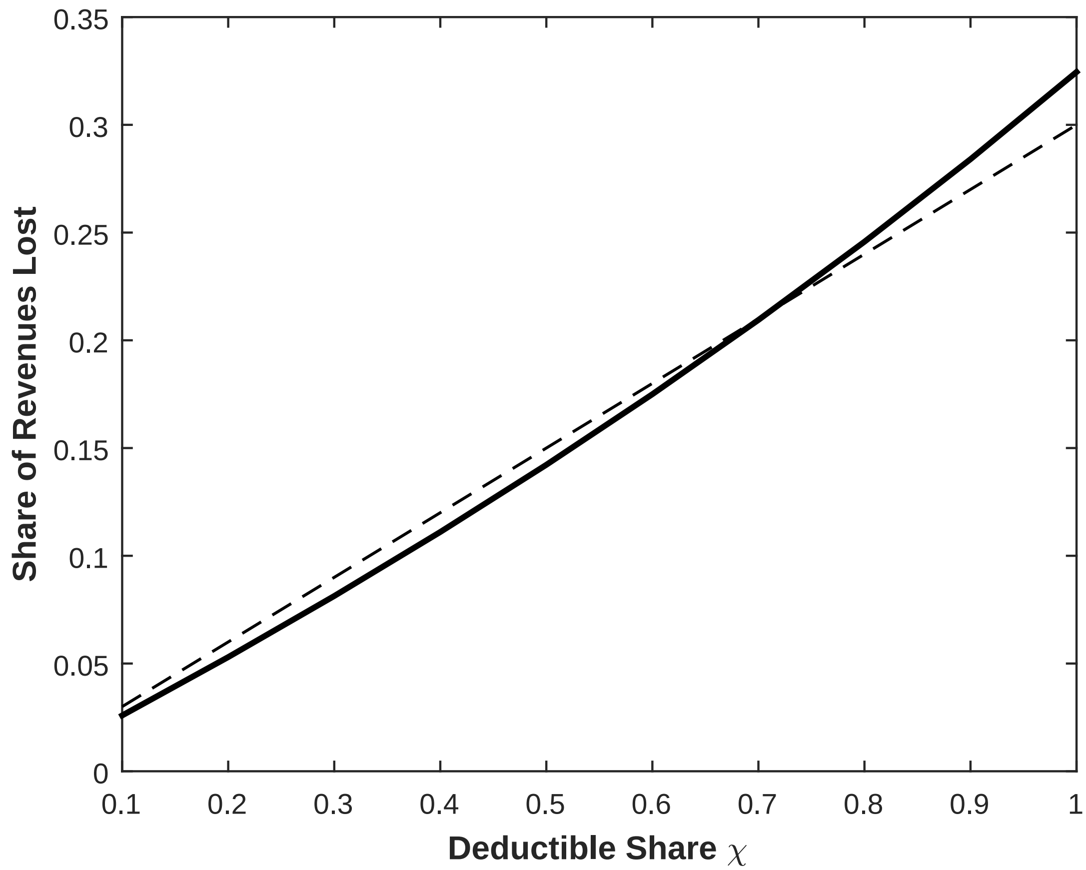

This figure shows the share of revenues lost due to a narrow tax base as a function of χ, the share of expenditure on exempt goods that is tax deductible. The figure is plotted for a value of $f = 0.7$, reflecting that 30% of goods are tax deductible. The solid line gives the actual share of revenues lost. The dashed line approximates revenues lost with a back-of-the-envelope calculation given by $(1 - f)\chi$.

Figure A.4: Magnitude of Tax Exemption and the Tax Base

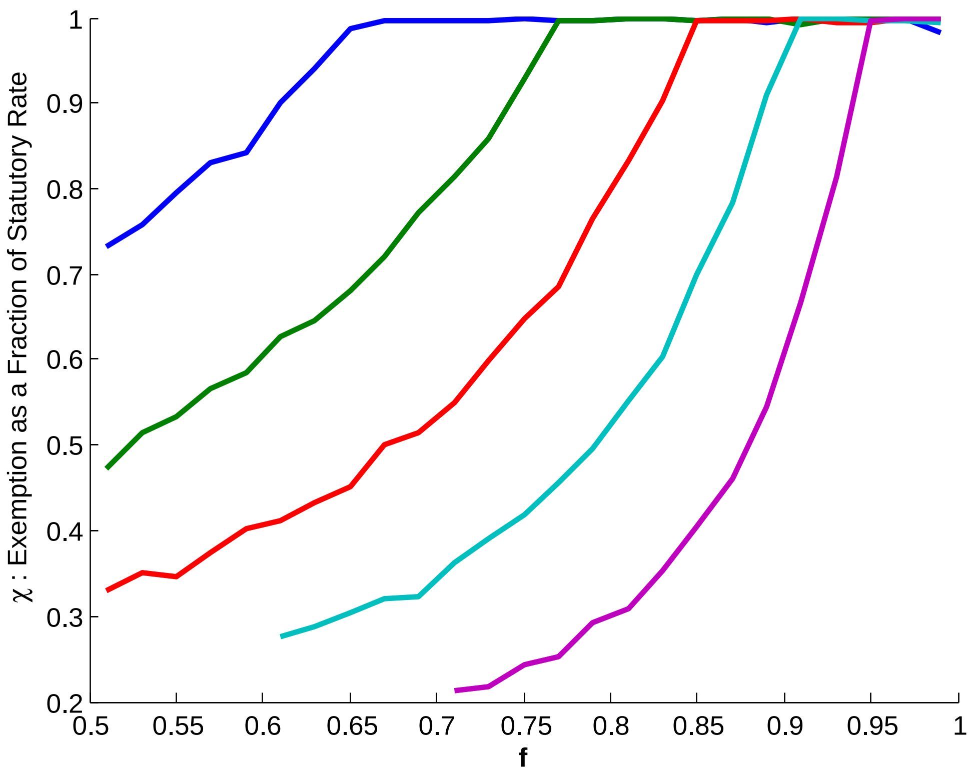

This figure shows the size of an exemption χ that would be chosen by a group of lobbyists of measure $1 - f$ as a function of the tax base f. χ is given as the fraction of expenditures on tax-deductible goods. Each curve is along a budget curve for a single value of g. The curves further to the right raise higher revenues. Above a critical threshold, lobbyists give themselves a full exemption. Below that threshold, lobbyists give themselves a smaller tax break as the tax base narrows.

Figure A.5: Utility of Exempt and Taxed with Equilibrium Partial Exemptions

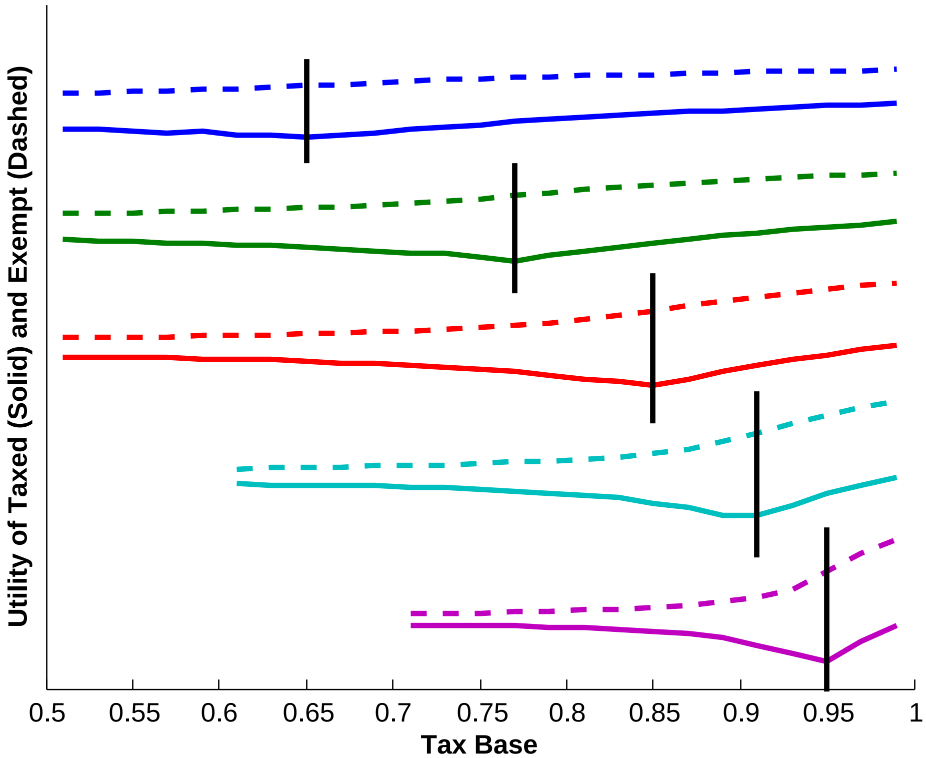

This figure shows the utility of citizens in equilibrium as a function of the tax base. Citizens whose firms have a tax exemption are given with a dashed curve, while the solid reflects taxed citizens. This utility results when a measure $1 - f$ jointly decide on the magnitude of a tax exemption, given in figure A.4. Horizontal lines give the inflection point, below which lobbyists give themselves only a partial exemption.

Figure A.6: The Benefit of Lobbying with Partial Exemptions

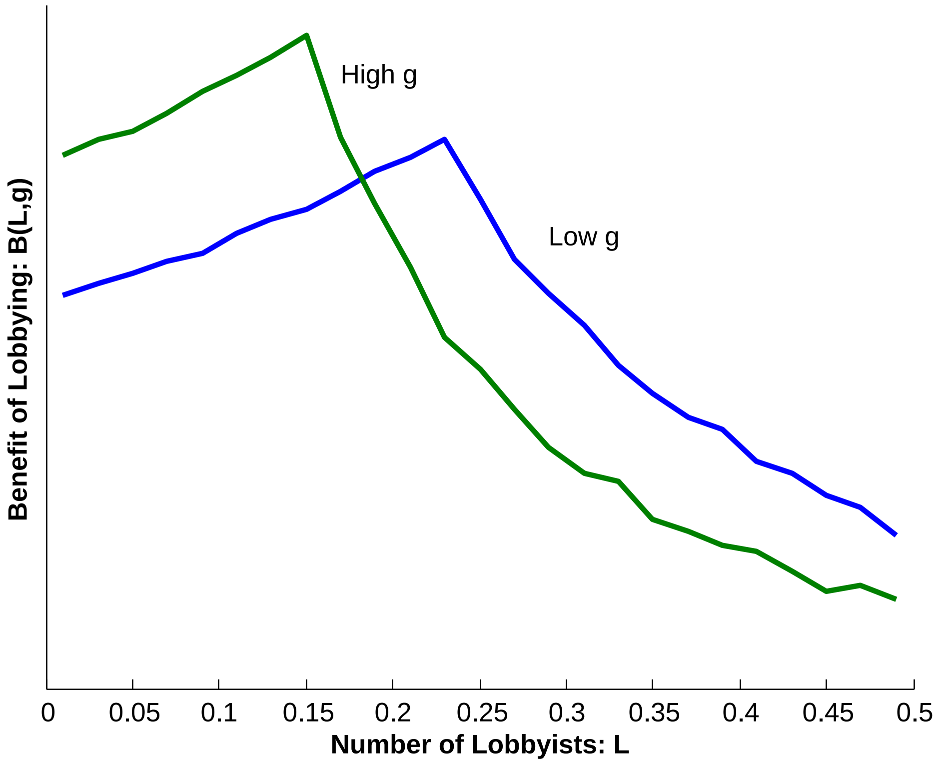

This figure shows the benefit of lobbying as a function of the number of lobbyists L. The benefit of lobbying is the difference between the utility of exempt and taxed citizens at a tax base of $1 - f$: the gap between the dashed and solid lines in figure A.5. The figure is plotted for two values of g.

Figure A.7: Example of Mixed-Strategy Equilibrium in Medium Voter Game

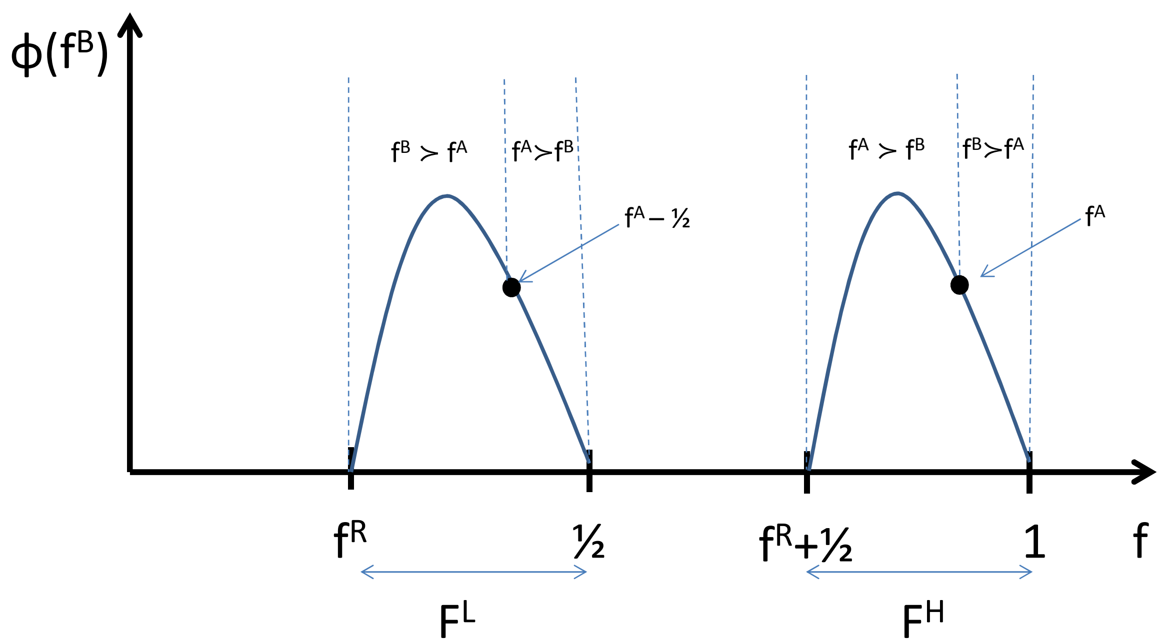

Figure A.8: Equilibrium Policy in Generalized Political Weights Model

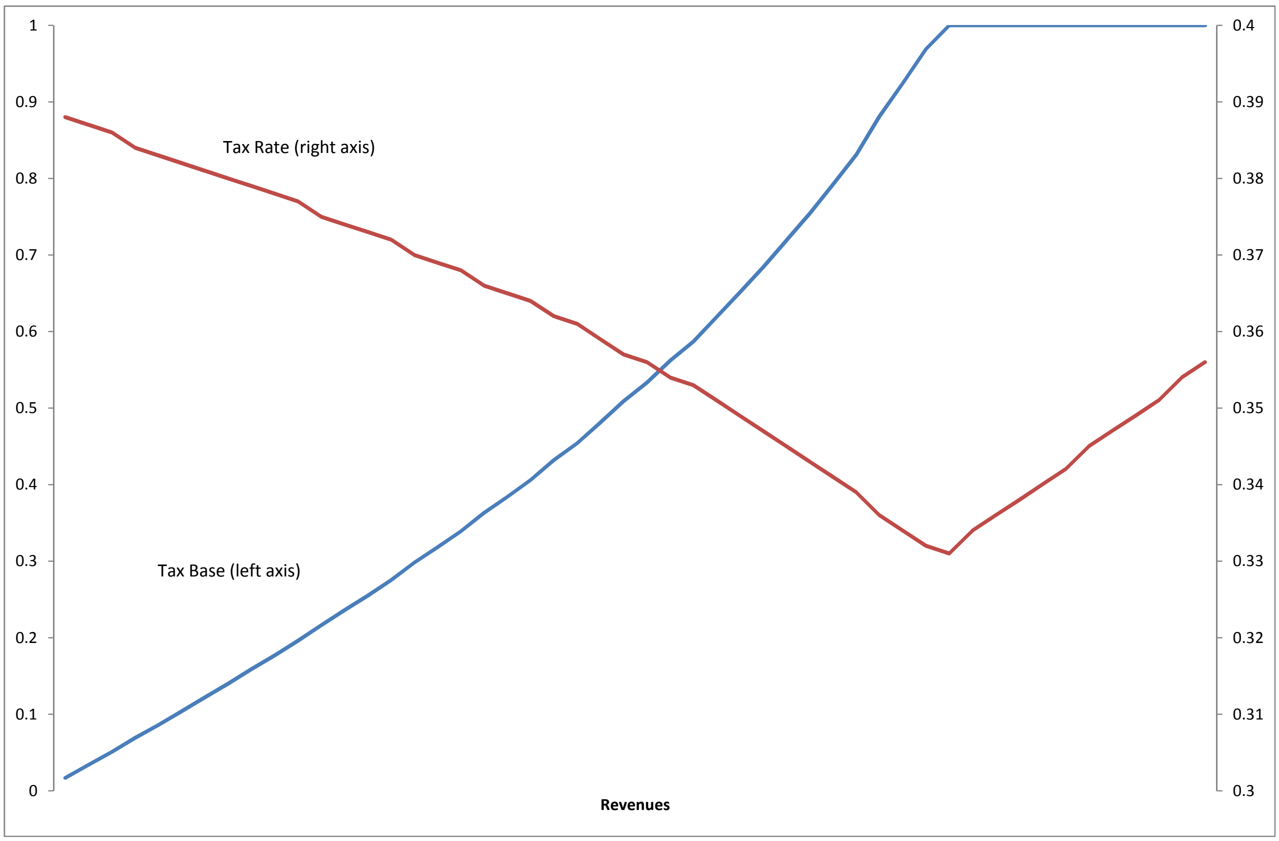

This figure shows the equilibrium tax base and tax rate as a function of g, in a general model where a social planner puts different weights on the utility of different citizens. The tax base expands and statutory rate declines as public good needs increase.