WORKING PAPERS

N° 1796

OCTOBER 2026

"The Fiscal Alibi: Hidden Spending Needs and Government Reputation"

Georgy Lukyanov and Hengrina Ly

Toulouse School of Economics

The Fiscal Alibi: Hidden Spending Needs and Government Reputation

Georgy Lukyanov* Hengrina Ly†

Abstract

A government may ask for a high tax because it faces a genuine expense. The same demand can also be made by a government that intends to keep the proceeds. We study this ambiguity in a two-period reputation model with privately observed spending needs and an endogenous tax base. The opportunist chooses between outright confiscation and levies that an honest government might impose. We characterize the three possible regimes through a single equilibrium equation. A mean-preserving spread of legitimate needs weakly raises the opportunist's lifetime value, with a strict increase precisely when the spread changes the upper tail relevant for mimicry. We then give a necessary and sufficient curvature condition for this ordering to extend to a general reputational continuation prize. Verification reduces the value of concealment, but also changes how the opportunist extracts. In the benchmark economy, greater auditing weakly lowers current citizen welfare, even before audit costs, while improving the future allocation through better information. The welfare case for verification consequently depends on the balance between these two effects. Examples yield no auditing, an optimum within a regime, and an optimum at the boundary between regimes.

Keywords: government reputation; hidden fiscal needs; taxation; convex order; audits; political agency.
JEL codes: D82; D83; H11; H21; H61.

1 Introduction

Suppose a government raises taxes to pay for an unexpected expense. Citizens observe the tax demand, but cannot readily establish how much the government needs to spend. A high tax is then consistent with two rather different explanations: the expense may be large, or the government may intend to keep some of the money. This ambiguity matters even if citizens have no difficulty observing how much they pay. What they cannot observe is whether the payment is justified.

*Toulouse School of Economics; georgy.lukyanov@tse-fr.eu.
†Université de Rouen Normandie, LERN; ENS Paris-Saclay, CEPS; hengrina.ly@univ-rouen.fr.

The possibility of a genuine expense gives an opportunistic government a reason that citizens cannot immediately reject. It can levy the same amount as an honest government facing a large bill. We refer to this possibility as a fiscal alibi.[1] The government has an incentive to preserve the ambiguity because its reputation affects future production and, hence, the amount it can subsequently extract. Its current levy must balance the revenue collected today against the reputation left for tomorrow.

This paper develops a model of that tradeoff. We ask how the distribution of legitimate spending needs affects opportunistic behavior, and whether verifying those needs improves citizen welfare. The two questions are related, but their answers need not have the same sign. Reducing the value of concealment can make a government more willing to confiscate current output. It is therefore necessary to distinguish the rents that a government obtains over its lifetime from the losses that its behavior imposes on citizens at a particular date.

There are two periods and two permanent government types. An honest government privately observes a spending requirement, collects exactly that amount, and finances the required service. An opportunistic government has no valued expenditure. It can collect an ordinary levy that resembles legitimate finance, or confiscate the entire tax base and reveal its type. Citizens supply labor before observing the levy. Their willingness to work depends on the probability that the government will leave the marginal product of their labor with them. Reputation has a continuation value because citizens make another labor-supply decision in the second period, when the honest government respects their property and the opportunist seizes it.

The distinction between an ordinary levy and confiscation is relevant here. An ordinary levy is a fixed assessment at an individual's margin; confiscation takes the individual's whole output. Thus, citizens retain the return to additional work whenever either type uses an ordinary levy. A government that conceals its type can sustain more current production than one that is expected to confiscate immediately. This feature allows us to isolate the effect of reputational cover. We return to proportional taxation in an extension, where the ordinary tax also distorts labor supply and some of the comparative statics change.

Our first result characterizes the equilibrium. An opportunist will not imitate every honest levy. A sufficiently small levy yields too little current revenue, even when it establishes honesty with certainty. Larger levies are attractive, but citizens must then attach a lower probability to honesty. In equilibrium, the loss of continuation value exactly offsets the increase in current revenue throughout the range that the opportunist uses. This indifference condition, together with Bayes' rule, determines the opportunist's likelihood relative to the honest levy distribution.

Integrating that likelihood gives the probability of an ordinary levy and therefore the current tax base. The entire equilibrium can then be recovered from one scalar equation. Depending on the parameters, the opportunist confiscates with certainty, mixes between ordinary levies and confiscation, or uses ordinary levies with certainty. We call the last

[1] There is no separate report of the spending need in the model. The opportunist chooses a levy, and citizens compare that action with the levies that an honest government could impose. The term *alibi* refers to this inference.

regime full mimicry. Full mimicry preserves the return to current labor, but does not eliminate extraction: every levy imposed by the opportunistic type is still diverted.

Two qualifications accompany this characterization. First, we assume that even the tax base supported by the honest type alone can finance the largest legitimate need. This keeps the comparison separate from insolvency and political replacement. Second, the uniqueness result concerns an explicitly defined class of regular perfect Bayesian equilibria.[2] These assumptions allow us to characterize the three regimes and their boundaries without adding a default convention or an announcement game.

The second result concerns fiscal risk. Holding the opportunist's value fixed, its relative likelihood is flat at zero over the range of small needs that it does not imitate. Above that range, the likelihood is increasing and strictly convex in the need. A mean-preserving spread therefore increases the total likelihood of mimicry at the original value. Equilibrium is restored by a weak increase in the opportunist's value. This argument applies across regime boundaries as well as within a regime.

The qualification to this result is economically useful. Additional dispersion has no effect if it only redistributes probability among needs that the opportunist would never imitate. What matters is the upper tail that can justify a sufficiently large levy. We give an exact condition for a strict increase in value in terms of this part of the distribution. The result concerns discounted opportunistic rents; it does not, by itself, establish a corresponding ordering of current diversion or citizen welfare.[3]

We next ask how much this conclusion depends on the continuation economy. For a general increasing reputational prize $C(\mu)$, where μ is the posterior probability of honesty, the relevant curvature condition is $2C'(\mu)+\mu C''(\mu) \geq 0$. This condition is sufficient for the distributional ordering throughout the regular, globally solvent class. It is also necessary for that ordering to hold for every local mean-preserving spread: if the condition fails, one can construct an equilibrium in which a small spread lowers opportunistic value. The terminal labor economy in our benchmark satisfies the condition. The proportional-tax extension shows why it need not hold when reputation must also sustain a distortionary tax base.

Finally, we introduce verification after an ordinary levy has been collected. An audit certifies an honest government and exposes an opportunist. It reduces the continuation reward from successful concealment and eventually makes the opportunist abandon mimicry. However, the adjustment begins before that point. When the opportunist still mimics with certainty, weaker future rewards lead it to collect more through current levies. When it mixes, verification raises the probability of outright confiscation and reduces current production.

[2]With continuously distributed levies, individual actions have probability zero. We obtain beliefs from local likelihood ratios rather than allowing an independent choice of belief at each such action. The resulting uniqueness claim is conditional on this restriction; it is not a uniqueness claim for every possible specification of off-path beliefs.

[3]A distribution with uniformly larger needs also raises opportunistic value. That comparison changes the mean. Our risk comparison holds the mean fixed, following the distinction in Rothschild and Stiglitz (1970).

These responses imply that, in the benchmark model, current citizen welfare weakly falls with auditing even before administrative costs. This conclusion depends on what the audit does: it reveals information after collection, without recovering the diverted funds or imposing a fine.[4] Citizens nevertheless benefit in the second period, when better information allows them to supply more labor under an honest government and avoid supplying it to an exposed opportunist. We derive both components of welfare. Neither dominates in general, and the examples include an optimal audit of zero, an optimum within the full-mimicry regime, and an optimum at the transition to partial mimicry.

1.1 Related literature

The reputational mechanism is related to Phelan (2006) and Lu (2013), in which government behavior and beliefs interact with private economic decisions. More generally, Fudenberg et al. (2022) and Luo and Wolitzky (2025) study reputation when actions, observations, and private information shape what behavior can be sustained. Our two-period structure is more restrictive than a recurrent reputation game. Its advantage for the present question is that the distribution of spending needs enters equilibrium through an explicit likelihood function. This makes it possible to compare entire distributions and to identify when the comparison reverses.

The hidden justification for a public expenditure also connects the model to political agency. Rogoff and Sibert (1988), Coate and Morris (1995), and Besley and Smart (2007) examine how information problems affect fiscal choices and political accountability. Grossman and Van Huyck (1988) distinguish excusable contingencies from opportunistic repudiation in sovereign borrowing. In our setting, the ambiguity arises within a tax demand: a levy that covers a large legitimate need is observationally indistinguishable from an opportunistic levy of the same size.

Related questions arise in work on fiscal opacity, budget rules, and volatile resources. Sloof et al. (2026) study fiscal opacity and debt ceilings, while Siemroth (2024) considers private spending needs and the allocation of funds within organizations. Cagé (2009) relates aid volatility to rent extraction under asymmetric information. We focus on the distributional comparison rather than an optimal fiscal rule. The main result identifies which changes in hidden needs raise reputational rents, and the verification analysis shows why reducing those rents need not improve the current allocation.

The information problem differs from that in Ablyatifov and Lukyanov (2026). There, uncertainty concerns whether an authorized fiscal mandate is delivered. Here, citizens observe the levy but do not know how much legitimate spending was needed in the first place. A large realized need can consequently provide cover for a government with nothing to finance. The distinction is central to the comparative statics, rather than an alternative interpretation of the same signal.

Section 2 sets out the environment and the belief restriction. Section 3 characterizes

[4] Recovery or a current penalty would introduce an additional deterrent. We omit these instruments to identify the effect of verification through reputation. This is also why the analysis does not amount to the optimal verification contract studied, for example, by Townsend (1979).

equilibrium. Section 4 studies changes in the need distribution and the role of continuation curvature. Section 5 derives the effects of verification on behavior and welfare, and Section 6 illustrates the results. The appendices contain the proofs, a discussion of solvency, and the proportional-tax extension.

2 Environment

2.1 Government types and spending needs

There are two periods. The government has a permanent type $\theta \in \{H, B\}$, and citizens begin with the prior

$$p = \Pr(\theta = H) \in (0, 1). \tag{2.1}$$

An honest government, type H, is committed to financing a legitimate public service. In period 1 it privately observes the amount G required to provide that service, levies $T = G$, and spends the proceeds accordingly. An opportunistic government, type B, has no valued expenditure and keeps what it collects. Citizens observe the levy, but neither the need nor the contemporaneous use of the money. Section 5 allows these facts to be verified after collection.

The distribution F of legitimate needs is common knowledge. It has a continuous density f, strictly positive on a compact interval $[\underline{g}, \overline{g}]$. Positive density will allow us to specify beliefs at every levy in the honest government's support. It does not require the opportunist to use that entire support.

There is a unit mass of atomistic citizens, each of whom produces one unit of output per unit of labor. A citizen's per-period utility is

$$c - \frac{\ell^{1+\gamma}}{1+\gamma}, \qquad \gamma > 0. \tag{2.2}$$

An ordinary levy is a fixed per-capita assessment. A citizen can therefore increase consumption by working more without increasing that assessment. Confiscation is different: it takes the citizen's entire output. Citizens choose labor before learning which levy will be imposed.[5]

We allow citizens to borrow against other wealth to meet a fixed assessment, so there is no individual cash-on-hand constraint. Aggregate feasibility nevertheless requires the legitimate expense to be no greater than aggregate output L. The benefit of honest spending is $B(G)$. For welfare comparisons we use $B(G) = G$, under which the benefit offsets the honest levy. A different benefit function adds $p\mathbb{E}[B(G) - G]$ to welfare and does not affect the policy comparisons below.

[5]The fixed assessment isolates the informational role of the levy. It is not equivalent to a proportional tax. With a proportional tax, an ordinary levy also lowers the marginal return to labor; Appendix C examines the resulting change in the argument.

2.2 Timing, production, and solvency

In period 1, Nature draws the government type and, if the type is honest, its spending need. Citizens then choose labor using their prior and the anticipated government strategy. After production, the honest government levies its need. The opportunist chooses either an ordinary levy $T \in [0, L)$ or confiscation of L. A demand for the whole output is defined to be confiscation, rather than a separate ordinary levy. Citizens observe the action and update their beliefs. In the verification extension, an audit follows an ordinary levy.

There is no new spending need in period 2. Citizens choose labor given the posterior μ that the government is honest. The honest government leaves their output with them, whereas the opportunist confiscates it. The expected return to a marginal unit of labor is consequently μ, and terminal labor is

$$L_2(\mu)^\gamma = \mu, \qquad L_2(\mu) = \mu^{1/\gamma}. \tag{2.3}$$

This output is also the opportunist's terminal seizure prize:

$$C(\mu) = \mu^{1/\gamma}. \tag{2.4}$$

Both citizens and the government discount terminal payoffs by $\beta \in (0, 1]$. The terminal period thus gives reputation a payoff derived from a labor-supply decision.[6]

Let m denote the probability that the opportunist chooses an ordinary levy. In period 1, citizens retain the marginal return to labor whenever the government is honest or the opportunist uses such a levy. Their common labor choice therefore satisfies

$$L^\gamma = p + (1 - p)m. \tag{2.5}$$

The smallest tax base consistent with this condition occurs at $m = 0$. Denote it by

$$r = p^{1/\gamma}. \tag{2.6}$$

It is sustained by the possibility that the government is honest, even when an opportunist is expected to confiscate with certainty. We impose the following restriction on legitimate needs.

Assumption 2.1. *The essential supremum of the legitimate need satisfies*

$$0 \leq \underline{g} < \overline{g} < r < 1. \tag{2.7}$$

Every honest need can then be financed from the smallest equilibrium base. In particular, an honest levy is strictly smaller than L and is never confused with the action defined as confiscation. The restriction is stronger than solvency along a particular equilibrium path. Its purpose is to make all three regimes, including the outcome under perfect verification, comparable within one economy. Appendix B explains what is needed when the upper bound on needs exceeds r.

[6] The absence of a second spending shock is an assumption about the continuation economy. It allows us to study the first-period information problem without specifying how a government defaults or is replaced after an unfinanceable need. It should not be read as a stationary model of recurrent public expenditure.

2.3 Beliefs and equilibrium

We next specify how citizens interpret an ordinary levy. This requires care because a single levy has probability zero under the continuous honest distribution. Let M be the opportunist's conditional measure over ordinary levies, with total mass $m \leq 1$, and let $I_\varepsilon(g) = (g - \varepsilon, g + \varepsilon)$. An assessment is *regular* if it satisfies the following conditions:

(a) The honest levy distribution is F. Whenever the denominator is positive, the posterior at an ordinary levy is obtained from the local likelihood ratio

$$\mu(g) = \lim_{\varepsilon \downarrow 0} \frac{pF(I_\varepsilon(g))}{pF(I_\varepsilon(g)) + (1-p)M(I_\varepsilon(g))}.$$

(b) This limit exists at every $g \in [\underline{g}, \overline{g}]$.

(c) Confiscation, any singular component of M, and ordinary levies outside the support of F receive posterior zero.

The first two conditions make beliefs consistent with the relative probabilities of nearby levies. The third treats actions that the honest type does not generate as evidence of opportunism. We will derive absolute continuity of the opportunist's equilibrium levy measure; it is not imposed on its strategy at the outset.

A regular perfect Bayesian equilibrium is a regular assessment in which citizens and the opportunist act sequentially rationally, while the honest type follows its commitment to levy and spend G. The uniqueness results below concern outcomes in this class.[7]

3 Equilibrium: betrayal and mimicry

Consider first the opportunist's choice after labor has been supplied. Confiscation yields L immediately and reveals the government's type, leaving no terminal tax base. An ordinary levy can yield less today while preserving a positive continuation payoff. Denote the opportunist's equilibrium lifetime value by V. On every levy that it uses with positive likelihood, the two components of its return must add up to this value:

$$g + \beta \mu(g)^{1/\gamma} = V. \tag{3.1}$$

The equilibrium ordinary-levy measure is absolutely continuous with respect to F, as established in the proof below. Write $h = dM/\,dF$ for its relative likelihood. This is a likelihood relative to the honest distribution, rather than a density with respect to the levy itself. Bayes' rule gives

$$\mu(g) = \frac{p}{p + (1-p)h(g)}. \tag{3.2}$$

[7]The local-likelihood requirement restricts beliefs at null actions. We do not prove that every sequence of fully mixed perturbations selects this requirement, and do not assert uniqueness when beliefs at such actions are left unrestricted.

Substitution into the indifference condition determines that relative likelihood:

$$h_V(g) = \frac{p}{1-p}\psi_V(g), \qquad \psi_V(g) = \left[\left(\frac{\beta}{V-g}\right)^{\gamma} - 1\right]_+, \tag{3.3}$$

Confiscation is always available, so $V \geq L \geq r > \overline{g}$. The denominator in (3.3) is therefore positive. The positive-part operator has a direct interpretation. If $g \leq V - \beta$, even a posterior of one cannot make the levy strictly preferable to the equilibrium return. The opportunist assigns it no positive density. We call this the frugal range. At larger levies, indifference requires progressively less favorable beliefs. The resulting posterior is

$$\mu_V(g) = \begin{cases} 1, & g \leq V - \beta, \\ ((V-g)/\beta)^{\gamma}, & g > V - \beta, \end{cases} \qquad g \in [\underline{g}, \overline{g}]. \tag{3.4}$$

Integrating the relative likelihood gives the probability that the opportunist uses an ordinary levy:

$$m(V) = \frac{p}{1-p} J_F(V), \qquad J_F(V) = \int \psi_V(g)\, dF(g). \tag{3.5}$$

Thus $J_F(V)$ aggregates the scope for mimicry at a given value. The corresponding ordinary-levy density with respect to g is $pf(g)\psi_V(g)/(1-p)$. The function ψ_V will also be the object used to compare need distributions. Its properties follow from the tradeoff already described: a larger current levy can compensate for a worse reputation, whereas a higher equilibrium value makes each particular levy less attractive.

Lemma 3.1. *For $V > \overline{g}$, $J_F(V)$ is continuous and non-increasing. It is strictly decreasing on $(\overline{g}, \overline{g} + \beta)$ and equals zero for $V \geq \overline{g} + \beta$. For each fixed V, $\psi_V(g)$ is non-decreasing and convex in g, flat below $V - \beta$, and strictly convex above it.*

It remains to make these choices consistent with labor supply and with probabilities adding up to one. If confiscation occurs with positive probability, the opportunist must be indifferent between it and its ordinary levies, so $V = L$. If it never confiscates, $m = 1$ and labor reaches its undistorted level $L = 1$, while the value of mimicry must be at least one. These observations give the following classification.

Theorem 3.2. *Under Assumption 2.1, there is a unique regular equilibrium outcome. Let $r = p^{1/\gamma}$.*

(i) ***Betrayal only.*** *If*

$$r \geq \overline{g} + \beta, \tag{3.6}$$

the opportunist confiscates surely and

$$V = L = r, \qquad m = 0. \tag{3.7}$$

(ii) ***Mimicry and betrayal.*** *If*

$$r < \overline{g} + \beta, \qquad p[1 + J_F(1)] < 1, \tag{3.8}$$

there is a unique $V \in (r, 1)$ *solving*

$$V^\gamma = p[1 + J_F(V)]. \tag{3.9}$$

The equilibrium probabilities and labor are

$$m = \frac{V^\gamma - p}{1 - p}, \qquad e = \frac{1 - V^\gamma}{1 - p}, \qquad L = V, \tag{3.10}$$

where e *is the opportunist's confiscation probability.*

(iii) ***Full mimicry.*** *If*

$$r < \overline{g} + \beta, \qquad p[1 + J_F(1)] \geq 1, \tag{3.11}$$

there is a unique $V \in [1, \overline{g} + \beta)$ *satisfying*

$$J_F(V) = \frac{1 - p}{p}. \tag{3.12}$$

The opportunist always mimics, $m = 1$*, confiscation is absent, and* $L = 1$*. Equality in* (3.11) *gives* $V = 1$*.*

Across all three regimes, V *is equivalently the unique root of*

$$\boxed{p[1 + J_F(V)] = \min\{V^\gamma, 1\}.} \tag{3.13}$$

The two regime transitions are continuous.

The scalar equation combines the government's incentives with the citizens' response. Its left-hand side is the marginal return to current labor implied by mimicry. Its right-hand side imposes indifference with confiscation whenever that action is used, and probability adding-up when it is not. The left-hand side is decreasing over the active range; the right-hand side increases up to one. Global solvency ensures that the crossing occurs where all legitimate levies remain feasible.

It is useful to distinguish full mimicry from honest behavior. More mimicry supports a larger tax base because it reduces the frequency of complete confiscation. Once $m = 1$, that base cannot increase further. Yet the opportunist continues to keep every ordinary levy it collects. A large base can therefore coexist with extraction that is concealed in legitimate-looking tax demands.

4 Spending risk and opportunistic rents

4.1 The global convex-order result

We now hold the prior, preferences, and continuation economy fixed, and change the distribution of legitimate needs. The distributions may have different compact supports, but each must satisfy Assumption 2.1. Write $F_1 \preceq_{cx} F_2$ when F_2 is a mean-preserving

spread of F_1: the means coincide and $\mathbb{E}_{F_2}\phi(G) \geq \mathbb{E}_{F_1}\phi(G)$ for every convex function ϕ (Rothschild and Stiglitz, 1970).

The comparison can be made in two steps. At the value generated by F_1, convexity of ψ_V implies that the spread weakly raises J_F. The left-hand side of the equilibrium equation consequently rises at the original solution. Since the difference between the two sides decreases through its crossing, the new equilibrium value cannot be lower. No assumption that the two economies remain in the same regime is needed.

Theorem 4.1. *Suppose $F_1 \preceq_{cx} F_2$, with each distribution satisfying global solvency. Then*

$$V(F_2) \geq V(F_1). \tag{4.1}$$

Across all regimes,

$$V(F_2) > V(F_1) \quad \Longleftrightarrow \quad J_{F_2}(V(F_1)) > J_{F_1}(V(F_1)). \tag{4.2}$$

In the interior regime, a strict value increase raises m and L and lowers e. In the full-mimicry regime, $m = L = 1$ already.

The strictness condition identifies which spending risks matter. At the original value, the kernel is zero below $c = V - \beta$. A spread entirely confined to that frugal range changes neither J_F nor the equilibrium. By contrast, additional risk can increase the opportunities for mimicry when it reaches levies above the cutoff.

To express this distinction directly in terms of tails, consider any fixed $V > \overline{g}^{\vee}$, where $\overline{g}^{\vee} = \max\{\overline{g}_1, \overline{g}_2\}$. For $c < \overline{g}^{\vee}$, the difference in mimicry functionals has the following representation:

$$\begin{aligned} J_{F_2}(V) - J_{F_1}(V) = \frac{\gamma}{\beta}\left\{\mathbb{E}_{F_2}(G-c)_+ - \mathbb{E}_{F_1}(G-c)_+\right\} \\ + \int_c^{\overline{g}^{\vee}} \frac{\gamma(\gamma+1)\beta^{\gamma}}{(V-t)^{\gamma+2}}\left\{\mathbb{E}_{F_2}(G-t)_+ - \mathbb{E}_{F_1}(G-t)_+\right\} dt. \end{aligned} \tag{4.3}$$

The quantity $\mathbb{E}(G-t)_+$ is the expected part of the need above a threshold t. Convex order weakly increases it at every threshold. Equation (4.3) weights these tail differences only at and above the mimicry cutoff, with strictly positive weights. If the cutoff is at or above both supports, the kernels and their difference are zero.

This representation also explains the limit of the welfare interpretation. The theorem orders the opportunist's discounted return, which includes both current revenue and its future seizure opportunity. It does not separately order current diversion. Nor does the improvement in current production when confiscation becomes less frequent establish a citizen-welfare gain.

A first-order stochastic increase in legitimate needs gives a separate comparison. Since ψ_V is increasing, such an increase also weakly raises value by the same crossing argument. It should be distinguished from the risk experiment, which keeps expected legitimate spending fixed.

4.2 A necessary and sufficient curvature condition

The power form of the terminal prize follows from the benchmark labor economy. To see which part of the risk result depends on that economy, replace the opportunist's terminal prize by a general function C. Keep the first-period fixed levy, labor supply, and global solvency assumptions unchanged. Suppose $C : [0,1] \to [0,1]$ is continuous, twice continuously differentiable on $(0,1)$, satisfies $C(0) = 0$, $C(1) = 1$, and has $C'(\mu) > 0$ in the interior.

The normalization retains a zero prize for an exposed opportunist and a unit prize at complete trust. Indifference on a mimicked levy now takes the form

$$g + \beta C(\mu) = V. \tag{4.4}$$

Inverting C gives the posterior needed to support any candidate return. Bayes' rule then converts that posterior into a relative likelihood. For $x \in (0,1]$, write

$$K_C(x) = \frac{1}{C^{-1}(x)} - 1. \tag{4.5}$$

The analogue of ψ_V, including the range in which the opportunist does not mimic, is

$$\kappa_{V,C}(g) = \begin{cases} 0, & g \le V - \beta, \\ K_C((V-g)/\beta), & g > V - \beta. \end{cases} \tag{4.6}$$

Its integral is

$$J_{F,C}(V) = \int \kappa_{V,C}(g)\, dF(g), \tag{4.7}$$

The piecewise definition keeps the inverse within its domain. The equilibrium equation remains (3.13), with $J_{F,C}$ replacing J_F. Monotonicity of C is enough for the classification and the unique crossing. The distributional comparison requires an additional curvature restriction.

Theorem 4.2. *The likelihood kernel generated by C is convex in the claimed need for every V if and only if*

$$\boxed{2C'(\mu) + \mu C''(\mu) \ge 0 \quad \text{for every } \mu \in (0,1).} \tag{4.8}$$

Under (4.8), Theorem 4.1 extends to C in every regular, globally solvent regime. Conversely, if (4.8) fails at some μ_0, there are parameters and a sufficiently small local mean-preserving spread around an active need for which the opportunist's equilibrium value strictly falls.

The necessity statement is about the scope of a general theorem. Failure of (4.8) does not imply that every spread lowers rents. It means that one can place an active equilibrium in a region of concavity and find a local spread that does so. The proof constructs this comparison with positive continuous densities and strict global solvency.

The benchmark $C(\mu) = \mu^{1/\gamma}$ satisfies the condition strictly for every $\gamma > 0$. A continuation economy with proportional taxation can instead have a fiscal-capacity threshold below which no positive financing base exists. Near that threshold, the relevant kernel can be concave. Appendix C develops this separate economy and gives explicit reversals. Because its continuation technology is not defined as a normalized positive financing branch over every posterior, the appendix is an illustration of the boundary, rather than a direct application of the theorem without further qualifications.

5 Verification and citizen welfare

5.1 Global equilibrium effects

Suppose citizens can verify the fiscal need and the use of funds after an ordinary levy has been collected. The audit takes place with a common probability $\alpha \in [0, 1]$, independently of the levy and the government type. An audited honest government is certified, while an audited opportunist is exposed. Conditional on the levy, the absence of an audit conveys no further information.

Verification changes terminal beliefs. It does not recover current revenue or impose a fine, and citizens anticipate its probability when they choose labor. For an opportunist that mimics, the continuation prize is received only when the levy escapes verification. Thus, for $\alpha < 1$, the equilibrium is obtained by replacing β in the government's incentive condition with

$$\widetilde{\beta} = \beta(1 - \alpha). \tag{5.1}$$

The citizens' discount factor remains β. This distinction will matter for welfare: $\widetilde{\beta}$ measures the opportunist's effective reward from concealment, not society's valuation of the future.

Under perfect verification, no ordinary levy can preserve the opportunist's reputation. It then prefers confiscation to any ordinary levy below L. The resulting base is r, which still finances every honest need by Assumption 2.1. Write $V(\alpha)$, $m(\alpha)$, and $L(\alpha)$ for the equilibrium value, mimicry probability, and labor.

Proposition 5.1. *Suppose the unaudited economy is active, so $r < \overline{g} + \beta$, and define*

$$\alpha_D = 1 - \frac{r - \overline{g}}{\beta} \in (0, 1). \tag{5.2}$$

Then:

(i) $V(\alpha)$ is continuous and non-increasing, and strictly decreasing while mimicry is active. The mimicry probability and labor are weakly decreasing, while outright confiscation is weakly increasing.

(ii) For $\alpha \geq \alpha_D$, the economy is in betrayal only: $V = L = r$ and $m = 0$.

(iii) If the unaudited economy has full mimicry, there is a unique $\alpha_F \in [0, \alpha_D)$ such that $V(\alpha_F) = 1$. Full mimicry holds below α_F, interior mimicry and betrayal hold on (α_F, α_D), and betrayal only holds above α_D. If the unaudited economy is already interior, the first region is absent.

Both transitions are continuous. Current concealed revenue need not be monotone in α.

The threshold α_D is the point at which even the largest honest levy, accompanied by the best possible unaudited reputation, is no better than confiscation. At and beyond that point, an ordinary levy already identifies the honest type. Additional audits provide no information that actions do not reveal.

Before that threshold, the adjustment depends on the initial regime. In full mimicry, the probability of an ordinary levy remains one, so labor remains at one. The opportunist instead changes the distribution of its levies to compensate for a smaller future reward. The welfare calculation below shows that expected current diversion rises on this branch. Once the economy enters partial mimicry, further verification also raises confiscation and reduces labor. A fall in lifetime value therefore need not take the form of less extraction today.

5.2 Welfare accounting

To evaluate auditing, we need to account for both periods and for the change in behavior. For $\alpha < 1$, let $\psi_{V,\alpha}$ be the kernel in (3.3) with $\widetilde{\beta}$ in place of β. Two averages will be useful:

$$D(V, \alpha) = \mathbb{E}\left[G\psi_{V,\alpha}(G)\right], \tag{5.3}$$

$$Q(V, \alpha) = \mathbb{E}\left[\min\left\{1, \frac{V - G}{\widetilde{\beta}}\right\}\right]. \tag{5.4}$$

The ex-ante amount diverted through ordinary levies is pD. To see the weighting, the opportunist's conditional density is $pf\psi/(1-p)$, and its ex-ante probability is $1-p$. The second average, Q, is expected terminal labor following an unaudited honest levy. In betrayal only, an ordinary levy reveals honesty; we set $D = 0$ and $Q = 1$, including at $\alpha = 1$.

Let $\mathcal{C}(\alpha)$ be the real resource cost of maintaining audit capacity, with $\mathcal{C}(0) = 0$. Citizens bear this cost through a separable lump-sum charge, so it does not alter labor incentives or the feasibility of legitimate spending.[8]

Using the normalization $B(G) = G$, citizen welfare is

$$\boxed{W(\alpha) = \frac{\gamma}{1+\gamma} L(\alpha)^{1+\gamma} - pD(V(\alpha), \alpha) + \frac{\beta p \gamma}{1+\gamma}\left[\alpha + (1-\alpha)Q(V(\alpha), \alpha)\right] - \mathcal{C}(\alpha).} \tag{5.5}$$

[8]This is an ex-ante capacity cost, rather than a fee incurred only when a particular levy is investigated. A cost of the latter kind would also depend on the equilibrium probability of an ordinary levy.

The first term is current output retained by citizens, net of labor costs and with honest spending credited at its benefit. The second subtracts ordinary levies kept by the opportunist. Confiscation is already reflected in the probability of retaining output and in equilibrium labor, so it must not be subtracted again.

The third term is discounted terminal surplus. An audited honest government induces the full-information labor choice, while an exposed opportunist induces no labor. Following an unaudited levy, labor is instead based on the posterior associated with that levy. This information effect can be made particularly transparent by writing $q = \widetilde{\beta} = \beta(1-\alpha)$ and expressing the terminal component as

$$W_2(\alpha) = \frac{\beta p \gamma}{1+\gamma} - \frac{p\gamma}{1+\gamma}\mathbb{E}\left[(G+q-V(q))_+\right]. \tag{5.6}$$

Here $V(q)$ is the same equilibrium value parameterized by the effective continuation weight. The positive-part term measures the loss relative to full information. On every active branch $V_q < 1$. Reducing q by increasing auditing therefore reduces that loss and raises terminal citizen welfare. This improvement must be compared with the current behavioral response.

Proposition 5.2. *Consider uniform audits under the benchmark continuation technology and suppose the unaudited economy is active.*

(i) Gross current citizen welfare is weakly decreasing in α in every regime. It falls strictly in the interior regime and, in full mimicry, whenever the active need distribution is non-degenerate. It is constant in betrayal only.

(ii) Gross intertemporal citizen welfare is continuous on $[0,1]$. If $\mathcal{C}$ is continuous, net citizen welfare is continuous and an optimal audit exists. If, additionally, $\mathcal{C}$ is non-decreasing, some optimum can be selected with $\alpha^* \leq \alpha_D$; if it is strictly increasing, no optimum exceeds α_D.

(iii) Even when $\mathcal{C} \equiv 0$, the sign of the intertemporal citizen-welfare effect is not fixed. The terminal information gain in (5.5) can outweigh the current loss, but need not do so.

The proposition separates a current loss from a future gain. While the opportunist still conceals its type, the prospect of preserving a future base restrains what it takes immediately. Verification weakens that restraint. Citizens gain later because they can condition production on better information about the government. There is no general reason for the second effect to dominate the first.

If the unaudited economy is already in betrayal, both actions reveal type and gross welfare is constant in audit intensity. A non-decreasing capacity cost then makes zero auditing optimal, though it need not be the only optimum if the cost is flat. More generally, auditing beyond α_D has no allocation or information benefit in this model. The relevant policy choice lies between no verification and the point where actions themselves fully separate the types.

6 Numerical illustrations

Two examples help distinguish the changes in rents from the welfare case for verification. Both use uniform needs and $\gamma = 1$, for which the mimicry functional can be integrated in closed form. The calculations locate equilibrium and regime boundaries continuously; the grid is used only to display the resulting curves.[9]

For the risk comparison, set $p = \beta = 0.5$. Compare $G \sim U[0.15, 0.35]$ with $G \sim U[0.05, 0.45]$. Expected spending is 0.25 in both economies, and the second distribution is a mean-preserving spread of the first. Both satisfy global solvency. Without audits, both equilibria involve partial mimicry, and the spread changes value and the mimicry probability as follows:

$$V : 0.64573 \longrightarrow 0.66736, \qquad m : 0.29146 \longrightarrow 0.33472. \tag{6.1}$$

The conditional probability of confiscation falls from 0.70854 to 0.66528. The larger value therefore coincides with a larger current base. The left panel of Figure 1 shows how auditing removes the additional cover. In the wider-need economy, even costless verification lowers intertemporal citizen welfare over the range where it changes behavior. Zero auditing is optimal.

For the policy example, set $p = 0.7$, $\beta = 1$, and $G \sim U[0.4, 0.6]$. The unaudited economy has full mimicry and $V = 1.20476$. With no audit cost, welfare is maximized at $\alpha_F = 0.29534$, where full mimicry gives way to partial mimicry. A cost of $\mathcal{C}(\alpha) = 0.1\alpha^2$ leaves the optimum at that boundary. With $\mathcal{C}(\alpha) = 0.2\alpha^2$, the optimum moves to $\alpha^* = 0.24888$, strictly within the full-mimicry regime. The right panel displays all three cases.

The positive audit optima do not eliminate concealment. At the lower optimum, the opportunist still uses an ordinary levy with certainty. At the regime boundary, confiscation is just becoming competitive. The welfare gain comes from better terminal information, despite the increase in current diversion.

7 Conclusion

The spending needs of an honest government affect the opportunities available to an opportunist. When citizens see the levy but cannot verify its justification, a government with nothing to finance can imitate the demands of a government facing a large expense. The distribution of those expenses then becomes a determinant of reputational rents.

The model identifies the relevant distributional comparison. A mean-preserving spread weakly increases opportunistic value because the likelihood required to support mimicry is convex in the legitimate need. The increase is strict only when the spread changes the upper tail used for mimicry. Risk among needs that are too small to be imitated has no

[9] The parameter values illustrate the mechanisms rather than estimate their size in a particular country. The replication files include the closed-form solver and a separate validator that reconstructs the examples by numerical integration.

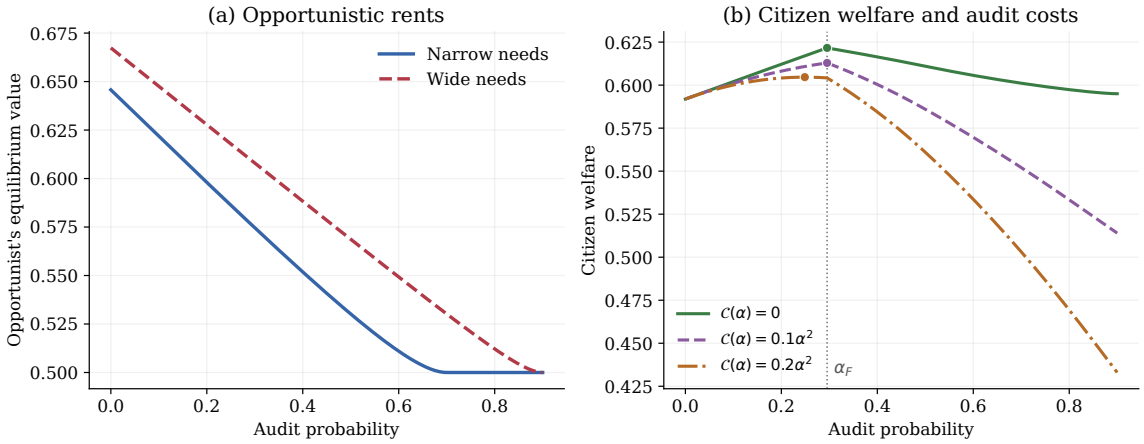


Figure 1: Spending risk and verification. Left: opportunistic value for the narrow and wide uniform distributions, with $p = \beta = 0.5$. Right: citizen welfare for $p = 0.7$, $\beta = 1$, and $G \sim U[0.4, 0.6]$, under three audit-cost functions. Dots mark the continuously located optima; the vertical line marks the full-to-partial-mimicry boundary α_F.

effect. The general continuation-prize result states the curvature condition behind this argument and shows how a local reversal can arise when it fails.

Verification reduces the future reward from concealment. Its current consequences are less favorable for citizens: the opportunist demands more through ordinary levies or becomes more willing to confiscate, and current welfare weakly falls. Better information improves the subsequent production decision. Whether auditing is worthwhile depends on that future gain relative to the current loss and the resource cost of verification.

These conclusions rely on the stated fiscal instruments and on a continuation economy in which reputation determines a future seizure opportunity. The proportional-tax extension demonstrates that changing the marginal tax incentive can change the distributional and audit responses. Recurrent spending needs, fiscal caps, or stabilization insurance would require further choices about financing, default, and political replacement. The present analysis leaves those choices open while identifying the reputational effect that such models would have to take into account.

A Proofs

Proof of Lemma 3.1. We first establish the shape of the likelihood kernel, and then integrate it. In the active range $g > V - \beta$, its derivatives are

$$\frac{\partial \psi_V(g)}{\partial g} = \gamma \beta^\gamma (V - g)^{-\gamma - 1} > 0, \qquad \frac{\partial^2 \psi_V(g)}{\partial g^2} = \gamma(\gamma + 1)\beta^\gamma (V - g)^{-\gamma - 2} > 0. \tag{A.1}$$

In the frugal range the kernel is zero. The two pieces agree at the cutoff, where the slope increases from zero on the left to γ/β on the right. The joined function is therefore non-decreasing and convex, with strict convexity on the active range.

For each need, the kernel is non-increasing in V. On a compact subset of $V > \overline{g}$, the denominator stays bounded away from zero, so dominated convergence gives continuity of its integral. If $V < \overline{g} + \beta$, there is positive probability above $V - \beta$, where the decrease in V is strict. The integral is consequently strictly decreasing there. If $V \geq \overline{g} + \beta$, no need is active and the integral is zero. □

Proof of Theorem 3.2. The proof has three steps. We first identify the form of any regular equilibrium strategy. We then solve the consistency condition for its value. Finally, we verify that the resulting strategy leaves no profitable deviation.

Under the belief restriction, a singular ordinary action receives posterior zero. If it is in the honest support, its payoff is at most $\overline{g} < r \leq L$; if it is outside that support, its payoff is $T < L$. Confiscation strictly dominates in either case. Ordinary actions outside the honest support are dominated for the same reason. Thus the equilibrium measure satisfies $M \ll F$. Since confiscation is always available,

$$V \geq L \geq r > \overline{g}. \tag{A.2}$$

At common Lebesgue points of the densities, the local-likelihood rule gives (3.2). Positive mimicry density on $g \leq V - \beta$ would give a posterior below one and a payoff strictly below V. Hence the density is zero almost everywhere on that range. Conversely, a positive-F-measure set with zero mimicry density above $V - \beta$ would contain a common Lebesgue point with posterior one. At that point, the opportunist could obtain $g + \beta > V$. Such a hole in the active range cannot occur in equilibrium. Indifference therefore implies $h = h_V$ almost everywhere, and

$$m = \frac{p}{1-p} J_F(V), \qquad L^\gamma = p + (1-p)m = p[1 + J_F(V)]. \tag{A.3}$$

If confiscation is used, $m < 1$ implies $L < 1$, and indifference with confiscation gives $V = L$. This yields (3.9). If confiscation is unused, the ordinary probabilities must add up to one, giving $m = L = 1$ and (3.12). The value must then satisfy $V \geq 1$ to deter confiscation. Together these conditions give (3.13).

Suppose first that $r \geq \overline{g} + \beta$. Even the largest honest levy with the best possible reputation pays no more than confiscation at the base r. Betrayal only is therefore an equilibrium, with the outcome in (3.7). No active equilibrium can have $V \geq r$ in this case, since its active set would be empty.

Now suppose $r < \overline{g} + \beta$. Define the difference between the two sides of the equilibrium equation on $[r, \overline{g} + \beta]$:

$$D(V) = p[1 + J_F(V)] - \min\{V^\gamma, 1\}. \tag{A.4}$$

At r, the active set has positive probability and $D(r) = pJ_F(r) > 0$. At the upper endpoint $J_F = 0$, while $\overline{g} + \beta > r$, so $D < 0$. Lemma 3.1 makes this difference strictly decreasing through its crossing. There is consequently one root. Its position below, at, or above one is determined by $p[1 + J_F(1)] - 1$. This gives the two active regimes and the probabilities stated in the theorem.

To verify the candidate, use the continuous density $h_V f$ and its local-likelihood posterior. Every mimicked levy yields V. Every unmimicked honest levy has $g \le V - \beta$ and hence yields at most V, even at posterior one. An off-support ordinary levy gives posterior zero and pays less than $L \le V$; confiscation pays $L \le V$. All deviation constraints are satisfied. Honest levies are feasible because $L \ge r > \overline{g}$. At the betrayal boundary $J_F(r) = 0$, and at the full-mimicry boundary $V = 1$. Substitution into the probability and labor formulas establishes continuity at both transitions. □

Proof of Theorem 4.1. Global solvency ensures that each equilibrium value exceeds both support bounds. On this common domain, convexity of the kernel and the definition of convex order imply

$$J_{F_2}(V) = \mathbb{E}_{F_2}\psi_V(G) \ge \mathbb{E}_{F_1}\psi_V(G) = J_{F_1}(V). \tag{A.5}$$

Evaluate this inequality at $V(F_1)$. The difference between the left- and right-hand sides of the equilibrium equation for F_2 is non-negative there. Since that difference crosses zero only once and decreases through the crossing, its root is weakly larger. It is strictly larger precisely when the inequality is strict at $V(F_1)$. This establishes both the ordering and its exact strictness condition, including comparisons across regimes. The labor and probability comparisons follow from (3.10) and the saturation $m = L = 1$ under full mimicry.

For the tail representation, take $V > \overline{g}^{\vee}$ and $c = V - \beta < \overline{g}^{\vee}$. The slope jump at c, together with the second derivative on the active range, gives

$$\psi_V(g) = \frac{\gamma}{\beta}(g - c)_+ + \int_c^{\overline{g}^{\vee}} \frac{\gamma(\gamma+1)\beta^{\gamma}}{(V-t)^{\gamma+2}}(g-t)_+ \, dt, \qquad c = V - \beta, \tag{A.6}$$

Taking expectations under each distribution and subtracting yields (4.3). If $c \ge \overline{g}^{\vee}$, both kernels vanish on the relevant supports, so their difference is zero. □

Proof of Theorem 4.2. We first derive the curvature condition and check that the active kernel joins its flat part correctly. Write $x = C(\mu)$. Differentiating the inverse in $K_C(x) = 1/\mu - 1$ gives

$$K_C''(x) = \frac{2C'(\mu) + \mu C''(\mu)}{\mu^3 [C'(\mu)]^3}. \tag{A.7}$$

The denominator is positive. Since $x = (V - g)/\beta$ is affine in g, the numerator also determines curvature with respect to the need. To check the join at $g = V - \beta$, let $A(\mu) = \mu^2 C'(\mu)$. Then

$$A'(\mu) = \mu[2C'(\mu) + \mu C''(\mu)], \qquad \frac{\partial K_C((V-g)/\beta)}{\partial g} = \frac{1}{\beta A(\mu)}. \tag{A.8}$$

Under (4.8), A is non-decreasing in μ. As the need rises, its supporting posterior falls, so the slope of the active kernel rises. At the cutoff, its one-sided limit is non-negative and therefore no smaller than the zero slope of the frugal range. The value is continuous

there. The condition is thus sufficient for convexity of the entire kernel, not only its active piece. Necessity follows from (A.7) at any interior posterior where the condition fails. Under convexity, the crossing argument in Theorem 4.1 applies to $J_{F,C}$.

For the converse, suppose the condition fails at μ_0. By continuity, the active kernel is strictly concave in a neighborhood of the corresponding need. We construct an interior equilibrium in that neighborhood. Choose $V \in (0,1)$, set $p = \mu_0 V^\gamma$, and thus obtain $r = \mu_0^{1/\gamma} V$. By taking V sufficiently small, we can choose $\beta \in ((V-r)/C(\mu_0), V/C(\mu_0))$ while maintaining $\beta \leq 1$. Set

$$g_0 = V - \beta C(\mu_0) \in (0, r). \tag{A.9}$$

With a point mass at g_0, the integral is $J_{F,C}(V) = 1/\mu_0 - 1$, and the equilibrium equation becomes $p[1 + J_{F,C}(V)] = V^\gamma$. The candidate is interior and strictly solvent.

It remains to implement the comparison within the regular density class. Approximate the point mass by a smooth positive density F_1^ε on a sufficiently small common support around g_0. Continuity of the unique crossing implies $V_1^\varepsilon \to V$. For small ε, the kernel at this equilibrium remains strictly concave near the central need. Spread a small central mass symmetrically within that neighborhood, leaving a common positive background density. The resulting distribution F_2^ε is a mean-preserving spread of F_1^ε, retains global solvency, and satisfies $J_{F_2^\varepsilon,C}(V_1^\varepsilon) < J_{F_1^\varepsilon,C}(V_1^\varepsilon)$. Its unique crossing is therefore at a strictly lower value. Smallness of the perturbation preserves the interior regime.

For the benchmark prize, direct differentiation gives

$$2C'(\mu) + \mu C''(\mu) = \frac{1}{\gamma}\left(1 + \frac{1}{\gamma}\right)\mu^{1/\gamma - 1} > 0. \tag{A.10}$$

□

Proof of Proposition 5.1. For an imperfect audit, the relevant government incentive is obtained by substituting $\widetilde{\beta} = \beta(1-\alpha)$. Holding V fixed, a higher $\widetilde{\beta}$ increases the kernel and its integral, strictly whenever some needs are active. The unique-crossing equation consequently gives a value that increases with $\widetilde{\beta}$ and decreases with auditing. The labor and probability comparisons follow from the formulas for the interior and full-mimicry regimes.

The betrayal threshold is reached when $r \geq \overline{g} + \widetilde{\beta}$. Solving this condition for α gives (5.2). If the initial equilibrium has full mimicry, the expression $p[1 + J_{F,\widetilde{\beta}}(1)]$ starts at or above one and falls continuously below one before α_D. Strict monotonicity while the active set is nonempty gives a unique crossing α_F. It may equal zero when the unaudited economy is already at the boundary. If the initial equilibrium is partial, this first transition is absent.

At the transitions, the values are respectively 1 and r, so both the value and the allocation formulas agree on either side. Perfect verification gives betrayal directly. Finally, $D(V, \alpha)$ changes both because the kernel depends on the effective weight and because equilibrium value adjusts. There is no general monotonicity of current concealed revenue

across the entire audit path; the full-mimicry and betrayal portions already exhibit different responses. □

Proof of Proposition 5.2. We begin with the accounting in (5.5). Expected retained period-1 output is $[p + (1-p)m]L = L^{1+\gamma}$. Subtracting labor cost gives $\gamma L^{1+\gamma}/(1+\gamma)$. Honest spending and its levy offset under $B(G) = G$; opportunistic ordinary levies have ex-ante measure $p\psi\, dF$, giving the deduction pD.

At an unaudited ordinary levy g, multiply expected terminal surplus by the probability density of observing that levy. Bayes' rule reduces the product to $p\, dF(g)\, \gamma\mu(g)^{1/\gamma}/(1+\gamma)$. With an audit, the full-information surplus $\gamma/(1+\gamma)$ is obtained only under the honest type. Integrating and discounting gives the terminal term in (5.5).

To compare behavior along the audit path, let $q = \widetilde{\beta}$. Define the following moments over the active need set:

$$A_j = \int (V-g)^j \left(\frac{q}{V-g}\right)^{\gamma} dF(g). \tag{A.11}$$

In partial mimicry, differentiate the equilibrium equation (3.9) and substitute into the derivative of current welfare. The result is

$$\frac{dW_1}{dq} = \frac{\gamma p}{q}\left[A_1 - \frac{pA_0^2}{V^{\gamma-1} + pA_{-1}}\right] > 0. \tag{A.12}$$

Cauchy–Schwarz gives $A_0^2 \le A_1 A_{-1}$, and the denominator is strictly larger than pA_{-1}. The displayed derivative is therefore positive. Since auditing lowers q, current welfare strictly falls. Differentiating the same equilibrium equation also yields

$$V_q = \frac{pA_0}{q(V^{\gamma-1} + pA_{-1})} < 1, \tag{A.13}$$

The active set satisfies $V - g < q$, implying $qA_{-1} > A_0$ and the strict upper bound on V_q.

In full mimicry, labor stays at one and $J_F(V) = (1-p)/p$. Differentiation of this adding-up condition and of D gives

$$\frac{dD}{dq} = \frac{\gamma}{q}\left(\frac{A_0^2}{A_{-1}} - A_1\right) \le 0. \tag{A.14}$$

Current welfare is a constant minus pD. Cauchy–Schwarz therefore implies that it falls weakly with auditing, and strictly if active needs are non-degenerate. On this branch $V_q = A_0/(qA_{-1}) < 1$. In betrayal only, current strategies and welfare do not change with the audit probability.

The identity $1 - Q = q^{-1}\mathbb{E}[(G + q - V)_+]$ gives (5.6). On an active branch, $V_q < 1$ means that the expected positive part rises with q. Thus terminal welfare rises with auditing, strictly while mimicry remains active. This establishes the opposite signs of the two gross welfare components.

At $V = 1$ and $V = r$, the equilibrium formulas and welfare components are continuous. A continuous capacity cost therefore makes net welfare continuous on the compact interval $[0, 1]$, ensuring existence of an optimum. For $\alpha \geq \alpha_D$, ordinary levies reveal honesty and confiscation reveals opportunism. Additional audits change neither allocation nor information. A non-decreasing cost allows selection of an optimum at or below α_D, and a strictly increasing cost excludes any larger optimum. Finally, the two costless examples in Section 6 establish that the intertemporal welfare response can have either sign. □

B Relaxing global solvency

The main solvency assumption gives a common feasible domain for every audit probability, including perfect verification. For a fixed imperfect audit, one can allow $\overline{g} \geq r$, provided $\overline{g} < 1$ is retained. The latter bound is necessary for strict financing feasibility because even full mimicry supplies no more than one unit of output.

A sufficient replacement condition is divergence of the mimicry integral at the upper support bound. With β replaced by the positive effective weight when auditing is present, require

$$\lim_{V \downarrow \overline{g}} J_F(V) = +\infty. \tag{B.1}$$

If the density is bounded away from zero near $\overline{g}$, the condition holds for $\gamma \geq 1$. More generally, if $f(g) \asymp (\overline{g} - g)^{\kappa}$, divergence occurs when $\kappa \leq \gamma - 1$.

The crossing argument can then start at $V \downarrow \overline{g}$, where its left-hand side diverges, instead of at r. Since $\overline{g} < 1$, it produces a strictly solvent active root: either partial mimicry with $L = V > \overline{g}$, or full mimicry with $L = 1 > \overline{g}$. When $\overline{g} \geq r$ and the audit is imperfect, betrayal only is not a separate regime under this divergence condition.

If the limit is finite, denote it by J_b. For $r < \overline{g} < 1$, a strictly solvent root exists if and only if

$$p(1 + J_b) > \overline{g}^{\gamma}. \tag{B.2}$$

This is the condition for a positive crossing gap at the lower endpoint $V = \overline{g}$. If it fails, the argument does not deliver a strictly solvent equilibrium of the specified form. Equality is a boundary case with no financing slack and requires separate treatment; a strict failure cannot support a root with every honest need below labor.

Even when the divergent-tail condition holds for every imperfect audit, perfect verification is different. It removes the continuation reward from mimicry altogether. If $\overline{g} > r$, the resulting betrayal base cannot finance some honest needs, and a default or outside-finance rule is needed. At $\overline{g} = r$, every need is only weakly financeable, and the top levy meets the action defined as confiscation. The strict inequality in Assumption 2.1 avoids both issues.

C Proportional taxation and fiscal capacity

This appendix changes the tax instrument as well as the continuation economy. Ordinary finance is now proportional to individual output, and the honest government also faces spending needs in the terminal period. The purpose is to show why the benchmark comparisons need not extend to that setting. We characterize a feasible upper financing branch and give counterexamples on it; we do not assert a global equilibrium classification for the rate economy.

Suppose terminal reputation is z, aggregate labor is X, and the honest government finances a need with mean $\mu_G > 0$ using the proportional rate G/X. The opportunist confiscates output. If every honest realization is feasible, the expected marginal retention rate is $z(1 - \mu_G/X)$. With labor curvature $\gamma \geq 1$, the symmetric labor condition is

$$X^{\gamma+1} = z(X - \mu_G). \tag{C.1}$$

A positive financing base requires $X > \mu_G$. Expressing the posterior as $z = X^{\gamma+1}/(X-\mu_G)$ shows that its minimum occurs at

$$X_* = \frac{\gamma+1}{\gamma}\mu_G, \qquad \underline{z} = \frac{(\gamma+1)^{\gamma+1}}{\gamma^\gamma}\mu_G^\gamma. \tag{C.2}$$

If $\underline{z} < 1$, each $z \in (\underline{z}, 1]$ has two positive financing roots, with a double root at $\underline{z}$ and none below it. If $\underline{z} > 1$, no posterior supports positive full financing; if $\underline{z} = 1$, only the double root at complete trust remains. These statements presume that every realized honest need can be financed, since otherwise the marginal tax wedge must incorporate the treatment of default.

We use the upper root. It is locally stable under the labor-response iteration $X_{\text{next}} = [z(1 - \mu_G/X)]^{1/\gamma}$: at a fixed point the derivative is $\mu_G/[\gamma(X - \mu_G)]$, which is below one on the upper branch and above one on the lower branch. This is a stability statement for that specified adjustment, rather than an additional equilibrium-selection theorem. For $\gamma = 1$, the threshold is $\underline{z} = 4\mu_G$ and the upper solution is

$$X_+(z) = \frac{z + \sqrt{z^2 - 4z\mu_G}}{2}. \tag{C.3}$$

Let $q = \beta(1 - \alpha) > 0$, and let X_1 be the upper solution at posterior one. At an active ordinary levy, indifference requires the continuation base $x = (V - g)/q$. Upper-branch inversion is valid only for $x \in [X_*, X_1]$. A need with $x \geq X_1$ is not mimicked and has zero likelihood; a candidate with $x < X_*$ cannot be supported on the selected branch. In particular, a continuous equilibrium supported on that branch requires $\overline{g} \leq V - qX_*$. Within the valid range, Bayes' rule gives

$$z = P(x) = \frac{x^{\gamma+1}}{x - \mu_G}, \qquad \phi(x) = \frac{1}{P(x)} - 1 = \frac{x - \mu_G}{x^{\gamma+1}} - 1. \tag{C.4}$$

Although the likelihood is increasing in the active need, it need not be convex. Differentiation with respect to g gives

$$\frac{\partial^2 \phi}{\partial g^2} = \frac{(\gamma+1)[\gamma x - (\gamma+2)\mu_G]}{q^2 x^{\gamma+3}}. \tag{C.5}$$

Near X_*, the numerator is negative. The part of the kernel associated with the smallest viable financing bases is therefore concave. A second relevant object is the retained-output contribution $\kappa_V(g) = (V-g)\phi((V-g)/q)$ in the first-period labor condition. Its curvature is

$$\kappa_V''(g) = \frac{\gamma}{q}\frac{(\gamma-1)x - (\gamma+1)\mu_G}{x^{\gamma+2}}. \tag{C.6}$$

For unit curvature, this expression is strictly negative throughout the active range. Thus neither the likelihood comparison nor the labor condition can be signed by importing convexity from the fixed-levy economy.

To obtain the first-period equation, consider an equilibrium in which the opportunist mixes with confiscation, so $V = L$. Under an honest government, the expected retained output is $V - \mu_G$. Under an opportunistic ordinary levy g, retained output is $V - g$, and the relative likelihood is $\phi((V-g)/q)$. Combining this marginal labor condition with Bayes' rule gives

$$V^{\gamma+1} = p\left[V - \mu_G + \int (V-g)\phi\left(\frac{V-g}{q}\right) dF(g)\right], \tag{C.7}$$

$$m = \frac{p}{1-p}\int \phi\left(\frac{V-g}{q}\right) dF(g), \tag{C.8}$$

Here m is the opportunist's conditional probability of an ordinary levy. The integrals apply on the valid upper-branch range, with zero likelihood at unmimicked needs. These equations are the partial-mimicry closure, not a formula for a full-mimicry regime in which value could exceed labor.

Consider the parameters

$$\gamma = 1, \qquad \mu_G = 0.2, \qquad q = 0.8, \qquad p = 0.75. \tag{C.9}$$

For a deterministic need $G = 0.2$, the admissible upper-branch mixing solution is $V_0 = 0.6000000$ and $m_0 = 0.6000000$. Now replace that need by the equal-probability spread $G \in \{0.16, 0.24\}$. On the same admissible branch, the scalar equation becomes

$$V^2 = 0.75\left[0.8 - 0.128\frac{V-0.2}{(V-0.2)^2 - 0.0016}\right] \tag{C.10}$$

Its unique admissible solution is

$$V_1 = 0.5957515, \qquad m_1 = 0.5987586. \tag{C.11}$$

The two continuation posteriors are 0.8607 and 0.8082, both above the financing threshold $4\mu_G = 0.8$. All honest needs are feasible. The spread therefore lowers both opportunistic

value and mimicry while remaining on the selected branch. Convolving both need distributions with the same sufficiently small continuous mean-zero noise preserves their convex-order ranking and, by continuity, the strict reversal. The displayed discrete example is not the source of the reversal.

Auditing can also have a different composition effect under proportional taxation. At a regular decreasing crossing, a higher audit rate locally reduces value, but the mimicry probability need not fall. For example, set $\gamma = 1$, $\mu_G = 0.2$, $p = 0.7$, $F = \delta_{0.2}$, and $\beta = 0.3$. With no audit, the admissible solution is $(V, m) = (0.3665074, 0.3557898)$. At $\alpha = 1/15$, the effective weight is $q = 0.28$, and $(V, m) = (0.3514951, 0.3851001)$. Even the probability of undetected mimicry increases: $(1 - \alpha)m = 0.3594267 > 0.3557898$.

The last comparison reduces rents and current labor, but raises mimicry and lowers confiscation. It illustrates the additional force introduced by the ordinary marginal tax wedge. Both the risk theorem and the audit-composition result in the main text consequently belong to the specified fixed-levy economy; extending them to proportional taxation requires further restrictions.

References

Ablyatifov, E. and G. Lukyanov (2026). Government reputation and fiscal capacity. Working paper, August 2026. arXiv:2509.03087v4.

Besley, T. and M. Smart (2007). Fiscal restraints and voter welfare. *Journal of Public Economics* 91(3–4), 755–773.

Cagé, J. (2009). Asymmetric information, rent extraction and aid efficiency. Paris School of Economics Working Paper 2009-45.

Coate, S. and S. Morris (1995). On the form of transfers to special interests. *Journal of Political Economy* 103(6), 1210–1235.

Fudenberg, D., Y. Gao, and H. Pei (2022). A reputation for honesty. *Journal of Economic Theory* 204, 105508.

Grossman, H. I. and J. B. Van Huyck (1988). Sovereign debt as a contingent claim: Excusable default, repudiation, and reputation. *American Economic Review* 78(5), 1088–1097.

Lu, Y. K. (2013). Optimal policy with credibility concerns. *Journal of Economic Theory* 148(5), 2007–2032.

Luo, D. and A. Wolitzky (2025). Marginal reputation. *Econometrica* 93(6), 2007–2042.

Phelan, C. (2006). Public trust and government betrayal. *Journal of Economic Theory* 130(1), 27–43.

Rogoff, K. and A. Sibert (1988). Elections and macroeconomic policy cycles. *Review of Economic Studies* 55(1), 1–16.

Rothschild, M. and J. E. Stiglitz (1970). Increasing risk: I. A definition. *Journal of Economic Theory* 2(3), 225–243.

Siemroth, C. (2024). Ending wasteful year-end spending: On optimal budget rules in organizations. *International Economic Review* 65(3), 1163–1188.

Sloof, R., R. Beetsma, and A. Steinweg (2026). Debt ceilings with fiscal intransparency and imperfect electoral accountability. *International Economic Review* 67(1), 97–125.

Townsend, R. M. (1979). Optimal contracts and competitive markets with costly state verification. *Journal of Economic Theory* 21(2), 265–293.