Discussion Paper Series

IZA DP No. 18695

May 2026

A Remedy to the Demand for Bad Policy

Riccardo Ghidoni
University of Bologna
and IZA@LISER

Giovanni Immordino
University of Naples Federico II
and CSEF

Paolo Roberti
University of Bozen-Bolzano

IZA@LISER NETWORK
ISSN: 2365-9793

IZA DP No. 18695 May 2026

A Remedy to the Demand for Bad Policy*

Abstract

Voters often oppose welfare-enhancing policies because they fail to anticipate how others will adjust their behavior. We show that policy design can align political support with efficiency even under biased beliefs. We study priority policies, a remedy whose off-equilibrium incentives make reform attractive even to biased voters. In a theory-guided experiment, participants vote between an inefficient status quo and a treatment-specific welfare-improving reform. A Pigouvian-like tax attracts only 27.5 percent support; adding priority incentives raises it to 41.7 percent, and a pure-priority policy raises it to 70 percent. Treatment effects are mostly driven by biased participants.

JEL classification

C92, D72

Keywords

reform, priority policy, voting, political failure, experiment

Corresponding author

Riccardo Ghidoni

riccardo.ghidoni@unibo.it

* *Acknowledgments*: We thank Nageeb Ali, Maria Bigoni, Federico Boffa, Marco Casari, Matteo Cervellati, Ernesto Dal Bó, Erik Eyster, Olle Folke, Edoardo Grillo, Marco Mantovani, Maximilian Mihm, Massimo Morelli, Antonio Nicoló, Melis Kartal, Giacomo Ponzetto, Eugenio Proto, Lucas Siga, Alistair Wilson, and seminar audiences at the Bocconi PERICLES group, Euregio workshop in Trento, Political Persuasion workshop in Venice, University Federico II, Alghero Workshop in Political Economy, University of Bologna, ESA World Meeting in Beijing, University Roma Tor Vergata, Populism workshop in Padova, Workshop on the Demand for Bad Policy in Bozen, University of Cagliari and Toulouse Business School. Paolo Roberti and Riccardo Ghidoni gratefully acknowledge funding through the Italian Ministry of University and Research (PRIN P2022C3XSS). We thank Marco Rosso for excellent research assistance. This study was approved by the University of Bologna's IRB (protocol 0217233, dated July 29, 2024).

1 Introduction

The persistence of inefficient ("bad") policies is a key societal problem. Institutional weaknesses, special-interest influence, and selection into politics are among the best-known supply-side explanations for this phenomenon.[1] Yet these factors do not fully explain why electorates themselves support inefficient policies. Voters often reject welfare-enhancing reforms, from urban congestion charges to carbon taxes, and reducing such opposition remains an open challenge for policy design.

Recent evidence suggests that voters' own beliefs and biases contribute to the demand for bad policy. Dal Bó et al. (2018) and Nunnari et al. (2024) show that a majority of individuals supports an inefficient status-quo policy even when an individually and socially optimal ("good") alternative is available. In their experiments, participants preferred to remain in a social dilemma game with no cooperation in equilibrium (the bad policy) rather than introduce a corrective tax that would yield substantial efficiency gains. This demand for bad policy stems from voters systematically underappreciating how a new policy affects others' behavior. These "pessimist" individuals fully internalize the new tax's cost on themselves, but underestimate the extent to which it induces others to cooperate, and therefore undervalue its benefits.

We take this demand-side failure as a policy-design problem. Rather than trying to correct beliefs directly, which can be difficult in practice, as evidence on tax perceptions suggests (Douenne and Fabre, 2022), we ask whether reform can be designed so that political support does not depend on correct equilibrium reasoning. We show that priority policies have this property. A priority policy modifies the status-quo social dilemma by awarding a prize to individuals who cooperate when others do not and imposing a punishment on those who defect when others cooperate. Like a standard corrective tax, it makes full cooperation the equilibrium. Unlike a standard tax, however, it makes the reform privately attractive off equilibrium even for voters who pessimistically expect others not to adjust. In this sense, priority incentives operate as contingent promises rather than as equilibrium transfers.[2] Once implemented, the reform induces mutual cooperation, so there is no lone cooperator to reward and no lone defector to punish.

The mechanism is easiest to see in our experimental setup. Consider two players playing a prisoner's dilemma (PD), our bad policy: They will not cooperate and the outcome is inefficient. They have the possibility to vote for a tax reform that, while lowering

[1] See, among others, Algan and Cahuc (2010); Knack and Keefer (1997); North (1990); Robinson and Acemoglu (2012); Blanes i Vidal et al. (2012); Luechinger and Moser (2014); Roberti (2019). On electoral accountability, see Besley (2005); Lee et al. (2004); on political selection, see Caselli and Morelli (2004); Gagliarducci and Nannicini (2013).

[2] The logic resembles deposit insurance in Diamond and Dybvig (1983): the instrument is valuable because of what it promises off the equilibrium path, but if it succeeds, the costly contingency is not triggered. Priority policies differ, however, in that they need not correct pessimistic beliefs; they make reform attractive even when voters continue to expect others not to adjust.

all payoffs, changes the incentives of the PD making cooperation the dominant strategy, yielding a higher payoff than the PD equilibrium. A rational voter, with accurate beliefs, will always support the tax. However, a pessimistic voter who expects the other player to keep defecting under the new good policy will not support it. Under a priority reform, by contrast, unilateral cooperation yields a higher payoff than the bad-policy equilibrium because the lone cooperator receives the priority prize—an incentive that is valuable off equilibrium but not paid once both players cooperate. The same pessimistic voter therefore supports the reform.[3] The mechanism works not because pessimists become accurate, but because the policy makes reform attractive even when they remain pessimistic. Thus, we contribute to the literature by taking biased beliefs as primitive and redesigning the reform so that it both promotes cooperation and wins political support.

Priority-like incentives are feasible in practice and can be implemented in terms of time or money. For example, in organ donation, where demand far exceeds supply, some countries give registered donors priority in receiving organs (Ashkenazi et al., 2015; Chandler et al., 2012; Kessler and Roth, 2012; Kim et al., 2021). A citizen who doubts that others will register as donors nevertheless has a strong incentive to register to improve their own waiting-list priority (the prize). If many others reason similarly and also register, the individual priority advantage vanishes, but the overall supply of organs increases, shortening waiting times for all. Similarly, in degressive rebate programs for green technologies, such as California's Solar Initiative (Hughes and Podolefsky, 2015), early adopters receive greater benefits than later adopters. The policy thus turns pessimism about others into a motive for early participation, producing momentum toward a socially beneficial outcome. Priority policies do not require intensive centralized monitoring. By linking the reward to an observable cooperative act, they induce cooperators to reveal themselves, as in donor registries; non-cooperators are then identified residually.

To test this policy-design principle, we conducted a laboratory experiment with 1,104 participants. In Part 1, participants play five rounds of the bad policy, i.e., the PD. At the start of Part 2, they vote between the status quo and a treatment-specific good-policy reform. The selected policy governs the next five rounds. This structure ensures that participants have experience under the bad policy, but must vote on the reform before experiencing it. In Part 3, they vote again and play one more round, allowing us to study learning from experience. The three baseline reforms are a corrective tax (*Tax* treatment) and two budget-balanced priority mechanisms: One layered on top of the corrective tax (*PriorityTax* treatment), and one without taxes (*Priority* treatment). In all these reforms, cooperation is the unique dominant-strategy equilibrium.[4] The experiment uses neutral

[3]The policy changes voters' incentives before adoption without requiring transfers along the equilibrium path; in the many-player extension, the same logic can be implemented with budget-balanced rewards and penalties that disappear at the efficient cooperative outcome.

[4]To benchmark behavior under each policy, we run counterfactual *Ctrl* conditions that mirror each treatment's sequence but assign the good or bad policy in Parts 2 and 3 exogenously, rather than by

language: Participants never see terms such as “tax,” “prize,” or “punishment”, which helps isolate incentive effects from framing. To test mechanisms, after voting but before the selected game is revealed, we elicit beliefs about cooperation under the good and the bad policies.

Consistent with our theoretical predictions, priority policies substantially reduce support for the bad policy. In Part 2, support for the bad policy declines from 72.5% in the Tax treatment to 58.3% in the PriorityTax treatment ($p = 0.031$). This 14.2 percentage-point decline is robust to the estimation strategy, a wide range of controls, and the exclusion of participants with poor understanding of the instructions. Yet demand for the bad policy remains substantial under PriorityTax. Under the pure Priority treatment, which removes taxes and implements larger budget-balanced priority incentives, bad-policy support falls to 30%, less than half its level under Tax ($p < 0.001$). Experience further reduces support for the bad policy in Tax and PriorityTax, but not in Priority, where support is already lowest across treatments.

The belief data are consistent with the proposed mechanism. Many participants are pessimistic about cooperation by others under the good policy, and those with more pessimistic beliefs are more likely to support the bad policy. This gradient, however, is weaker under PriorityTax and weaker still under Priority than under Tax, consistent with priority policies reducing opposition among pessimistic voters. Belief data also help explain why Priority is more effective than PriorityTax. Many supporters of the bad policy are not only pessimistic about the reform; they are also overoptimistic about cooperation under the bad policy. For these voters, the tax component in PriorityTax compresses the payoff gains from switching, so the priority prize cannot compensate for a status quo they already perceive as favorable. Removing taxes, as in the Priority treatment, widens the payoff differential and delivers substantially higher support for the good policy. Additional analyses and treatments suggest that neither status-quo bias nor participants’ failure to anticipate how their own behavior would change under the reform are the primary drivers of our findings.

Our findings identify a design principle for reducing demand for bad policy: Policies should remain privately attractive even for voters who reason pessimistically about others’ adjustment. Conceptually, this principle is reminiscent of obviously strategy-proof design (Li, 2017): The reform is structured so that voters need not forecast others’ behavior to see why supporting it is in their own interest. Experimentally, we isolate this logic in a simple, transparent 2×2 setting and show theoretically that it extends to many-player environments while preserving cooperation and budget balance. Our paper thus complements supply-side explanations of policy failure, information-based remedies, and demand-side accounts in which opposition to reform is fully rational (Ali et al., 2025), by showing how policy design itself can reduce politically inefficient choices.

voting.

The rest of the paper proceeds as follows. Section 2 reviews demand-side explanations for the persistence of inefficient policies and situates our contribution within this literature. Section 3 introduces the priority policy, derives the conditions under which it restores cooperation, and extends the analysis to many players. Section 4 describes the experimental design and derives theory-guided predictions. Section 5 presents the results on policy support, the mechanisms underlying it, and the role of learning. Section 6 concludes.

2 Related literature

Demand-side explanations for the persistence of bad policy tend to emphasize distortions in voters' own assessments of reform. In Fernandez and Rodrik (1991), voters oppose reform because they face individual uncertainty about whether they will gain or lose, a friction that would disappear if individual outcomes were known in advance. Ali et al. (2025) show that distributional concerns can generate rational opposition to efficient reforms when voters interpret others' support as a negative signal about a policy's consequences. In Glaeser and Ponzetto (2017), voters' systematic inference errors that overattribute people's actions to innate characteristics distort policy choices and political selection. Closest to our work, Dal Bó et al. (2018) identify a channel in which voters who underappreciate how reform changes others' behavior oppose welfare-improving policies, because the direct costs of reform are salient but the indirect benefits through behavioral adjustment are not.

Subsequent work reinforces the importance of beliefs about others' responses. Nunnari et al. (2024) show that support for efficient reforms increases with cognitive ability and, crucially, with beliefs about other citizens' ability to understand the reform and adjust their behavior. More broadly, a growing empirical literature documents widespread misperceptions about what others think and do, and shows that correcting such beliefs can change behavior (see, e.g., Bigoni et al., 2019; Bursztyn et al., 2020; Bursztyn and Yang, 2022). These contributions indicate that misperceptions about others are pervasive and consequential.[5]

While our paper sits close to Dal Bó et al. (2018) and Nunnari et al. (2024), it shifts the question from diagnosis to policy design: Can policy itself be structured so that political support is robust to voters' failure to anticipate others' equilibrium responses? We identify and test a behaviorally robust principle of reform design—that policies should remain privately attractive even under pessimistic beliefs—and show that priority sat-

[5] Biased beliefs about self-competence can also guide inferior collective choices (Kartal and Tyran, 2022). Arteaga et al. (2022) document that overoptimism and incorrect beliefs can lead to suboptimal school allocations across students. Furthermore, a large literature shows that voters are flooded with misinformation, making belief formation about equilibrium behavior in new policies difficult to achieve (Angelucci and Prat, 2024; Acemoglu et al., 2024).

isfies this principle while standard corrective taxation does not.[6] A recent paper by Dreyfuss (2026) complements our work by developing a general theoretical framework of equilibrium neglect, and providing survey evidence consistent with our experimental results.

Finally, a related literature studies priority rules in organ allocation as incentive schemes to increase donation, showing that donor-priority can raise registration rates, attract political support, and improve welfare (Chandler et al., 2012; Kessler and Roth, 2012; Kim et al., 2021). Our approach is complementary. Rather than studying whether priority affects behavior or support within a specific allocation system, we ask whether the priority *principle* can overcome demand-side opposition to welfare-improving reforms more generally. In this sense, our work also relates to the experimental literature on endogenous institution choice, which shows that democratically selected enforcement mechanisms can sustain cooperation (see, e.g., Dal Bó et al., 2010; Sutter et al., 2010). That literature typically asks whether democratic choice improves compliance; we ask what determines whether voters choose welfare-improving institutions in the first place, and how policy design can ensure that they do.

3 Theoretical background

This section develops the theory. Starting from a two-player PD as the bad status quo, we compare three good policies: a tax scheme, a tax–plus–priority hybrid, and a pure priority rule. By granting a verifiable benefit to those who cooperate when others defect and by penalizing those who defect when others cooperate, priority policies make cooperation a best reply even under pessimistic beliefs; this, in turn, should increase support for priority relative to the tax scheme. We characterize the parameter ranges under which the policy is budget-balanced and extend the logic to many players, preserving both budget balance and the cooperative equilibrium.

3.1 A priority policy

Consider a status quo in which individuals do not cooperate because the social benefit generated by their cooperative effort is fully appropriated by defectors. The canonical representation is the PD in Table 1a: Cooperation (C) costs $c > 0$ and produces benefit

[6] Recent work shows that resistance to taxation is driven not only by material incidence but also by beliefs, salience, and policy framing. In a representative French survey after the Yellow Vests protests, Douenne and Fabre (2022) show that carbon-tax opposition is strongly linked to pessimistic beliefs about private losses, regressivity, and policy effectiveness. More generally, Kroft et al. (2024) demonstrate that consumers respond more to price changes than to equivalent tax changes, implying that tax salience materially shapes incidence and welfare. Finally, Folke (2014) shows how political parties avoid competing in tax policy, focusing on secondary policies.

Table 1: Prisoner's Dilemma versus Tax Game

(a) Prisoner's Dilemma

	C	D
C	$b-c$, $b-c$	$-c$, b
D	b, $-c$	0, 0

(b) Tax Game

	C	D
C	$b-c-t_C$, $b-c-t_C$	$-c-t_C$, $b-t_D$
D	$b-t_D$, $-c-t_C$	$-t_D$, $-t_D$

$b > c$ for the other player. Defection (D) is therefore dominant and (D, D) is the unique Nash equilibrium, despite (C, C) being Pareto superior.

A standard remedy for social dilemmas is to tax the socially harmful action so that cooperation becomes individually optimal. However, such Pigouvian-like interventions may fail politically. Dal Bó et al. (2018) study support for a policy that imposes a small tax t_C on C and a larger tax t_D on D. Under $b > t_D > t_C + c$, the policy transforms the PD into the Tax Game (TG) in Table 1b,[7] where C is dominant and the unique dominant-strategy equilibrium is (C, C), yielding payoffs $b - c - t_C > 0$ (rather than 0 in the PD). Thus, a rational player who expects others to be rational would prefer the TG to the PD, and if a majority shared this view, the tax would be politically supported. In the experiment, however, many participants oppose this welfare-improving policy and prefer the PD, consistent with misjudging how incentives affect behavior.

Dal Bó et al. (2018) formalize this idea in a simple framework: Support for the PD over the TG arises when voters underestimate how others will adjust their behavior under the new policy. Rather than holding fully rational expectations, voters may have beliefs about both their own actions and others' actions under the two possible policy scenarios, PD and TG. If voters underestimate the other player's decision to cooperate in the TG (and/or overestimate the decision to cooperate in the PD), they will choose the PD over the TG, leading to a lower individual payoff and collective welfare.

We propose a remedy based on a priority policy designed to remain attractive even when voters underestimate equilibrium effects. Specifically, we augment the PD with (i) a priority prize P awarded to a player who cooperates when the other defects, and (ii) a punishment p imposed on a player who defects when the other cooperates. Think, for instance, of organ donation: Registered donors move up the waiting list (gain P), while non-registrants move down (loss p). If all players cooperate (i.e., they register for organ donation), no prize is given, but they all enjoy the social benefit of full cooperation, and individual defection is punished by p. The resulting Priority Policy Game (PPG) is shown in Table 2.

Cooperation is dominant in the PPG if both the prize P and the punishment p are greater than the cooperation cost c, i.e., $P > c$ and $p > c$. We also impose budget balance $P \leq p$, to limit the scope of priority policies to those that are self-financed.

[7] Dal Bó et al. (2018) call this game the Harmony Game.

Table 2: Priority Policy Game

	C	D
C	$b-c$, $b-c$	$-c+\boldsymbol{P}$, $b-\boldsymbol{p}$
D	$b-\boldsymbol{p}$, $-c+\boldsymbol{P}$	0, 0

Table 3: Priority Policy Game with Taxes

	C	D
C	$b-c-t_C$, $b-c-t_C$	$-c-t_C+\boldsymbol{P}$, $b-t_D-\boldsymbol{p}$
D	$b-t_D-\boldsymbol{p}$, $-c-t_C+\boldsymbol{P}$	$-t_D$, $-t_D$

Hence, the set of feasible policy parameters is $c < P \leq p$. The payoff table should be interpreted as a reduced-form representation of the incentive effects of priority. In a monetary implementation, the punishment paid by unilateral defectors can finance the prize received by unilateral cooperators; if $p > P$, the residual can be rebated or retained by the implementing institution. In the experiment, we set $P = p$, so no residual surplus arises.[8] As we show in the next section, unlike a Pigouvian-like tax, the priority policy can command political support and move society away from inefficient outcomes.

3.2 A simple framework

For comparability with the tax scheme, this section compares voters' incentives in the PD and in a modified PPG that incorporates the TG's tax structure, which we call the PPGT.[9] The voter holds belief α^B that she will cooperate under the bad policy and belief β^B that the other player does. Similarly, α^G and β^G denote her beliefs about her own and the other player's cooperation under the good policy.

Let us compare players' preferences over the PD and the PPGT. Consider a rational player who anticipates that the other player is also rational. In the PPGT, the unique dominant-strategy equilibrium is (C, C), so she would hold beliefs $\alpha^G = \beta^G = 1$. In the PD, by contrast, the unique dominant-strategy equilibrium is (D, D), implying beliefs $\alpha^B = \beta^B = 0$. As in the comparison between the TG and the PD, the rational player prefers the PPGT to the PD.

Consider now players who have non-equilibrium beliefs in the two games. They will vote for the PPGT if

$$\beta^B b - \alpha^B c \leq \alpha^G \beta^G (b-c-t_C) + \alpha^G (1-\beta^G)(P-t_C-c) + (1-\alpha^G)\beta^G (b-t_D-p) + (1-\alpha^G)(1-\beta^G)(-t_D)\,. \quad (1)$$

[8]In non-monetary applications, such as queue priority, the same logic is implemented through timing rather than through literal monetary transfers.

[9]The tax t_C can be set to 0 without affecting the subsequent results. In the TG, t_D must be sufficiently large for (C, C) to be the (unique) dominant-strategy equilibrium. In the PPG, either t_D or the punishment p must be sufficiently large for (C, C) to be the (unique) dominant-strategy equilibrium, so t_D and p are substitute incentives.

For ease of exposition, assume the player holds sufficiently low beliefs α^B about her own cooperation in the PD (the status quo) and β^B about the other player's cooperation. Such beliefs may arise from experience: after playing the PD for multiple rounds, behavior converges to the dominant-strategy outcome (see, e.g., Ghidoni et al., 2019), and beliefs adapt accordingly.

Assumption 1. *The player holds equilibrium beliefs in the PD:* $\alpha^B = \beta^B = 0$.

Given Assumption 1, we define three player types based on their beliefs about cooperation in the PPGT (or any good policy). *Accurate* types hold $\alpha^G = 1$ and $\beta^G = 1$; *pessimist* types hold $\alpha^G = 1$ and $\beta^G = 0$; and *inertial* types hold $\alpha^G = 0$ and $\beta^G = 0$. Accurate types correctly anticipate that both they and the other player will cooperate in the PPGT. Pessimists anticipate that they will cooperate in the PPGT but expect the other player to defect. Inertial types expect defection by both players in the PPGT and, in particular, fail to anticipate their own behavioral change across games.[10]

Consider a pessimist type. Unlike in the TG, this player, who contemplates cooperating under the PPGT, has an incentive to vote for the latter game over the PD because she expects to receive the prize P, especially when she believes the other player will defect under the PPGT. Proposition 1 formalizes this intuition.

Proposition 1. *A voter has a preference for the Priority Policy Game with Taxes over the Prisoner's Dilemma if the expected priority prize $\alpha^G(1-\beta^G)P$ is sufficiently large, i.e., she is sufficiently pessimist and the prize is sufficiently high, or she holds beliefs α^G and β^G sufficiently high, i.e., she is sufficiently accurate.*

Proof. A voter has a preference for the Priority Policy Game over the Prisoner's Dilemma if and only if her perceived gains from moving to the Priority Policy Game are positive:

$$\beta^B b - \alpha^B c \leq \alpha^G\beta^G(b-c-t_C) + \alpha^G(1-\beta^G)(P-t_C-c) + (1-\alpha^G)\beta^G(b-t_D-p) + (1-\alpha^G)(1-\beta^G)(-t_D).$$

When the voter believes both players to play the Nash equilibrium in both games ($\alpha^B = 0, \beta^B = 0, \alpha^G = 1, \beta^G = 1$), the previous inequality is satisfied, because $b > t_C + c$, while prize P and punishment p become irrelevant. Now consider $P = 0$. If the voter is sufficiently pessimistic about cooperation in PPGT, $\beta^G < \frac{\beta^B b - \alpha^B c + (1-\alpha^G)t_D + \alpha^G(t_C+c)}{b-p+\alpha^G p}$, the voter chooses the Prisoner's Dilemma. If instead $P > 0$, for $\beta^G < \frac{\beta^B b - \alpha^B c + (1-\alpha^G)t_D + \alpha^G(t_C+c)}{b-p+\alpha^G p}$, the voter chooses the PPGT if and only if the prize P multiplied by $\alpha^G(1-\beta^G)$ is sufficiently large:

$$\alpha^G(1-\beta^G)P > \alpha^G t_C + (1-\alpha^G)t_D + (\alpha^G - \alpha^B)c + \beta^G(1-\alpha^G)p + (\beta^B - \beta^G)b.$$

[10]The biased beliefs of pessimist and inertial types could be microfounded by assuming a positive correlation between beliefs in the PD and in the PPGT: these voters perceive their experience in the PD as carrying over to the PPGT (see Section 5.5).

□

First, observe the contrast in voters' behavior when PPGT, rather than TG, is presented as the good policy option. For the same beliefs, high α^G and low β^G, the player votes for the PD over the TG.[11] By contrast, when the option is the PPGT, the player prefers it to the PD provided the expected priority prize is sufficiently large. Second, if both players share these beliefs, they both vote for the PPGT and then choose to cooperate, resolving the social dilemma.

Turning to the gain from cooperation, we ask what is the minimum prize P that induces a pessimist voter to choose the PPGT over the PD? Consider a purely pessimist type, with $\alpha^G = 1$ and $\beta^G = 0$, and maintain Assumption 1. For this voter, the prize must offset the cost of cooperation and the tax on cooperation, i.e., $P > c + t_C$. Combined with the other parameter restrictions for the PPGT, this yields $c + t_C < P \leq p$. These inqualities will guide our choice of P in the experimental treatment arm with PPGT.

Notice that we can restate condition (1) for choosing the PPGT over the PD in terms of β^G:

$$\beta^G \geq \overline{\beta}^G := \frac{\beta^B b - \alpha^B c - \alpha^G\big(P - t_C - c\big) + (1-\alpha^G)t_D}{\alpha^G(b-P) + (1-\alpha^G)(b-p)}, \tag{2}$$

if $p \leq P < b$.[12] Given Assumption 1 (i.e., $\alpha^B = \beta^B = 0$), a purely pessimist type (with $\beta^G = 0$ and $\alpha^G = 1$) votes the PPGT because the threshold $\overline{\beta}^G$ becomes negative. If we relax Assumption 1 and allow the player to place a positive probability on the other player cooperating in the PD (i.e., $\beta^B > 0$), then a pessimist type may vote for the PD; this occurs whenever $\overline{\beta}^G > 0$. Empirically, the distribution of β^B is not fully concentrated at 0 (see Section 5 and Dal Bó et al. (2018)), so when Assumption 1 fails we obtain the following corollary.

Corollary 1. *If Assumption 1 is not satisfied, the set of beliefs β^G, for which a player votes for the PPGT, becomes larger if tax t_C or t_D decreases.*

Therefore, we expect support for a priority policy to increase as taxes are reduced or eliminated. The intuition is straightforward. If a voter assigns even a small probability to the event (D, C) in the PD (which yields payoff b), then a prize exceeding $c + t_C$ may not be sufficient to secure support for the PPGT. Moreover, consistent with Proposition 1, the threshold $\overline{\beta}^G$ also decreases in the prize P. When taxes are completely eliminated, PPGT becomes PPG, the pure priority game.

In Appendix A we investigate two relevant extensions of the theory: the n-player version of the theory with linear surplus, and the n-player version of the theory with

[11] See Proposition 1 in Dal Bó et al. (2018).

[12] If $b < P \leq p$, $\beta^G \leq -\overline{\beta}^G$: when prize and punishment become larger than b, then players hope that the counterpart is not cooperating, to access the large prize. Independent on the size of P and p, increasing taxes reduces the set of beliefs β^G for which a player votes for the PPGT.

decreasing marginal surplus. In both these extensions players, when assigned a game by the voting procedure, play the PD, TG, or PPG (with no taxes), with $n-1$ other players. We show that the results of the two-player version completely carry over to these extensions. Interestingly, the theory with decreasing marginal surplus answers a relevant policy concern relating to the adoption of priority policies: when the population of individuals is large, and many individuals already cooperate, can a prize for cooperation be sufficient to incentivize cooperation also in the remaining group of defectors? The answer is yes, because the efficient number of cooperators is such that the benefit added to the total surplus by the last cooperator must be larger than or equal to the individual cost of cooperation.[13] Consider the case where there is only one remaining defector: the prize given to her can always be as large as the full increase in total surplus implied by her cooperation, making her cooperation individually rational. Clearly, this is not true when individual cooperation implies an increase in total surplus lower than her cost of cooperation, but this is a case in which the number of cooperators is larger than the efficient number.

In the next section, we explain how we bring this environment to the laboratory to study whether actual behavior matches equilibrium behavior, whether subjects underappreciate the indirect effect due to the adjustment of others, whether this affects subjects' preferences over games, and whether our priority policy can mitigate this result.

4 The experiment

This section presents the experimental design and implementation. We first describe the general experimental setup (Section 4.1) and the treatment conditions (Section 4.2). We then state the testable predictions implied by our theoretical framework, focusing on the key treatment comparisons (Section 4.3). Finally, we report implementation details (Section 4.4).

4.1 Structure of the session

The structure of the experimental session is identical across all main treatments. Each session consists of 11 rounds of interaction. Participants are first randomly assigned to matching groups of six and play five rounds in Part 1. They are then reshuffled into new matching groups of six for five additional rounds in Part 2, and reshuffled once more to play a single round in Part 3. Within each matching group, participants are re-paired

[13] In a large population with decreasing marginal surplus, it can be optimal to have a group of defectors. The same result would be obtained, if the cost of cooperation is heterogenous across individuals: there would be a threshold on the cost of cooperation such that all individuals with a cost larger than the threshold should defect.

each round, ensuring that no participant encounters the same counterpart more than once within a given part.

All pairs within a matching group play the same game, either the good or the bad policy, in every round. After each round, participants are informed about their counterpart's action and their payoff. In the baseline treatments, participants begin under the bad policy in Part 1. At the beginning of Parts 2 and 3, all members of a matching group submit a choice of which game to play in that part, selecting between the bad and the good policy. The choice of one randomly selected group member (i.e., the random dictator) is then implemented for all rounds of that part. The lack of experience with the good policy, relative to the experience accumulated under the policy played in Part 1, mimics real-world settings in which voters are asked to express support for a reform *before* experiencing its effects. Participants are informed about which policy is implemented in each part, but neither the voting decisions of others nor the identity of the random dictator are disclosed, including to the random dictator.

At the beginning of Parts 2 and 3, we also elicit beliefs about cooperation under both the good and the bad policy. Specifically, after participants report their policy preference but before they learn which policy will be implemented in their group, they are asked to state the probability that two randomly selected participants from previous sessions cooperate under each policy. Belief elicitation in Part 2 is incentivized using the mechanism proposed by Karni (2009), which is robust to risk preferences. Following Dal Bó et al. (2018), participants are asked to estimate the behavior of participants from previous sessions to minimize the impact of belief elicitation on decision-making.[14] The same beliefs are elicited in Part 3, but accuracy is not incentivized to avoid potential hedging concerns.

The session concludes with a brief survey collecting sociodemographic information, including age, gender, highest educational qualification, and field of specialization. The survey also includes a simple guessing-game task commonly used to measure strategic reasoning (see, e.g., Nagel, 1995). Finally, the survey contains a free-form response section in which participants can elaborate on their policy choices in Parts 2 and 3.

4.2 Treatments

While the bad policy is identical across all treatments and corresponds to the PD, the good policy varies by treatment. We label our baseline treatments after the corresponding good policy. In the Tax treatment, the good policy is the Tax Game (TG) introduced in

[14] For each treatment, we first ran a smaller session in which beliefs were elicited but not incentivized. Behavior observed in Part 2 of these sessions is then used to score belief accuracy in subsequent sessions. In particular, we rely on treatment-specific control sessions, of which participants are informed, in which the Part 2 policy is assigned by a coin flip (see Section 4.2 for additional details). This ensures that reported beliefs reflect expectations about behavior unaffected by endogenous policy selection.

Figure 1: Experimental games across treatments

Bad policy

	C	D
C	$9, 9$	$3, 11$
D	$11, 3$	$5, 5$

Good policies

(a) Tax

	C	D
C	$8, 8$	$2, 7$
D	$7, 2$	$1, 1$

(b) PriorityTax

	C	D
C	$8, 8$	$7, 2$
D	$2, 7$	$1, 1$

(c) Priority

	C	D
C	$9, 9$	$11, 3$
D	$3, 11$	$5, 5$

Notes: C = Cooperate, D = Defect. The bad-policy game is identical across treatments and corresponds to a Prisoner's Dilemma; the good-policy game varies by treatment.

Section 3.1. In the PriorityTax treatment, the good policy is the Priority Policy Game with Taxes (PPGT), which combines the tax structure of the Tax Game with a priority prize and punishment. Finally, in the Priority treatment, the good policy is the Priority Policy Game (PPG), in which taxes are eliminated.

The payoffs associated with the good policy in each treatment are shown in Figure 1. In all treatments, we set $b = 6$ and $c = 2$, and add a constant of 5 to every payoff. This normalization makes all experimental payoffs positive and implies that the payoff matrix under the good policy in the Tax treatment is identical to that used by Dal Bó et al. (2018), facilitating a direct comparison with their results. The PriorityTax treatment combines the tax mechanism with a priority bonus and punishment, providing an intermediate benchmark for assessing how the priority policy performs relative to the tax mechanism studied in previous work. We choose its parameters to match the Tax treatment in terms of both total surplus and the Gini coefficient for each outcome profile. Thus, taxes on cooperation and defection are the same in the Tax and PriorityTax treatments, with $t_C = 1$ and $t_D = 4$. In PriorityTax, we additionally set the prize P and punishment p equal to 5. Finally, in the Priority treatment, we remove taxes altogether by setting $t_C = t_D = 0$ and introduce a budget-balanced priority policy with $P = p = 8$, chosen to yield outcomes comparable to those under the bad policy in terms of total surplus and the Gini coefficient.

To pin down behavioral mechanisms, we implement two additional treatments that are variants of PriorityTax. The first treatment, referred to as RevPriorityTax, is designed to rule out status quo bias as a potential driver of demand for the bad policy. In this treatment, participants begin Part 1 under the good policy. The second treatment, referred to as PriorityTaxPlus, is designed to examine the role of inertial players. In this treatment, payoffs are identical to those in PriorityTax, except that the mutual-defection payoff is 5 rather than 1.

Finally, for all treatments described above, we run corresponding counterfactual control conditions, labeled Ctrl. In these conditions, the sequence and type of games mirror those in the respective treatment, but the game played in Parts 2 and 3 is determined by a (virtual) coin toss rather than by a random dictator. These controls provide a clean counterfactual benchmark for behavior in Parts 2 and 3 when policy selection is exogenously assigned.

4.3 Predictions

Under the assumption of rationality, players would never prefer the bad policy to the good one in any treatment, since dominant-strategy equilibrium payoffs under all good policies strictly exceed those under the bad policy. Nevertheless, support for the good policy in the Tax treatment may be undermined if participants underestimate its equilibrium effects. In particular, if they anticipate limited behavioral adjustment following the reform by others or by themselves, they may opt for the bad policy.

Specifically, in the Tax treatment, a preference for the bad policy can arise if participants expect post-reform play to converge either to mutual defection, yielding a payoff of 1 to each player, or to a CD outcome, yielding a payoff of 2 to the player who (correctly) cooperates. In these cases, the expected payoffs are lower than the status-quo payoff of mutual defection in the PD, which is 5.

We test the prediction that participants facing a priority-policy reform are less likely to support the bad policy (see Proposition 1) by comparing voting behavior in the Tax treatment with that in the PriorityTax and Priority treatments. In these two treatments, participants who expect themselves to cooperate but anticipate defection by their partners should favor the good policy over the bad one. Indeed, conditional on cooperating while the partner defects, the good policy yields a payoff of 7 in PriorityTax and 11 in Priority, compared to the status-quo payoff of 5 under the bad policy, due to the presence of a priority bonus.[15] This leads to the following prediction:

Prediction 1. *Support for the bad policy decreases in the PriorityTax and Priority treatments, relative to the Tax treatment.*

Corollary 1 shows that support for the bad policy should decrease when taxes t_C and t_D are eliminated from PriorityTax, provided that a non-negligible fraction of participants assigns positive probability to cooperation by others under the bad policy.

Prediction 2. *Support for the bad policy decreases in the Priority treatment relative to the PriorityTax treatment.*

[15] This logic does not apply to inertial participants who expect both themselves and their partner to defect; we return to this case in Prediction 4 and in the PriorityTaxPlus treatment.

Corollary 1 also shows that the belief threshold about others' cooperation under the good policy, above which participants vote for the good policy, decreases when taxes are eliminated, as in the Priority Policy Game. A second change in moving from the PPGT to the PPG is that $P = p$ rises above b. When the priority prize is sufficiently large, participants who are pessimistic about the counterpart's cooperation under the good policy should be especially likely to support it, because they expect to receive the prize.

Prediction 3. *Increased support for the good policy in the Priority treatment relative to the PriorityTax treatment is concentrated among participants with lower beliefs about others' cooperation in the good policy.*

According to the theoretical framework, a preference for the bad policy would persist under the PriorityTax and Priority treatments only if participants also expect to defect themselves (i.e., inertial types). However, Dal Bó et al. (2018) show empirically that most participants anticipate greater behavioral attrition from others than from themselves. The PriorityTaxPlus treatment, which raises the DD payoff under the good policy from 1 to 5 relative to PriorityTax, is designed to shed further light on this mechanism. In the PriorityTaxPlus treatment, the bad policy and the good policy yield identical payoffs under mutual defection, reducing the expected cost of adopting the good policy for inertial types and making other payoff comparisons more salient.

Prediction 4. *Support for the bad policy decreases in the PriorityTaxPlus treatment relative to the PriorityTax treatment.*

While outcomes under the good policy in the Tax and PriorityTax treatments share the same Gini coefficient, they differ in terms of advantageous and disadvantageous inequality off equilibrium. Social preferences à la Fehr and Schmidt (1999) and Charness and Rabin (2002) combined with equilibrium reasoning, however, cannot drive our treatment effect, since equilibrium allocations are identical across the two treatments.

4.4 Procedures

The experiment was preregistered on OSF (`https://osf.io/85b7h`), programmed in oTree (Chen et al., 2016) and conducted by Playstudies between November 2024 and November 2025 at the LINEEX laboratory in Valencia.[16] Each treatment included 120

[16]The oTree code contained an unintended error in the implementation of the random-dictator rule: in some cases, the policy implemented for a matching group corresponded to the choice of a participant outside that group. The error affected all treatments. Importantly, in every matching group, the implemented policy was also chosen by at least one member of that group. Because participants observed only the implemented policy, and not the identity or vote of the selected dictator, each realized assignment was observationally equivalent to one that could have arisen under the intended rule, with the dictator being a group member who made the same policy choice. The error therefore did not change participants' incentives, information, payoff table, or feedback relative to such a valid implementation.

participants, except CtrlPriority, which included 24. Participants' total earnings consisted of the sum of payoffs from the eleven rounds, incentive payments for belief accuracy, the payoff from the guessing game, and a show-up fee of €5. Mean total earnings were approximately €12. Sessions lasted 45–60 minutes.

Instructions for each part were read aloud and displayed on participants' screens at the beginning of the corresponding part, and remained accessible throughout the experiment. The instructions used neutral language: terms such as taxes, priority, punishment, or prize were never mentioned. A sample is provided in Appendix C. At the beginning of Parts 1 and 2, participants were required to pass a five-question comprehension quiz before proceeding to the decision-making phase; they could attempt the quiz as many times as necessary without penalty. Communication among participants was not allowed.

5 Results

This section presents the experimental results. We focus primarily on treatment comparisons in Part 2. We first show that moving from the bad policy to any of the good policies substantially increases cooperation and efficiency (Section 5.1). We then estimate how PriorityTax and Priority affect participants' support for the bad policy relative to Tax (Section 5.2). Next, we examine the role of pessimistic beliefs about the good policy (Section 5.3), the role of inertial participants (Section 5.4), and overoptimism about the bad policy (Section 5.5). We also assess whether status-quo bias contributes to the demand for bad policy (Section 5.6). Finally, we investigate learning dynamics between Parts 2 and 3 (Section 5.7).

5.1 Good policies induce high cooperation

Figure 2 plots event-study estimates comparing changes in cooperation rates between Parts 1 and 2 for groups that move to the good policy with those that remain under the bad policy. Estimates in Part 1 are close to zero and display no systematic pattern across rounds, lending support to the identification assumption of parallel trends.[17] Moving to the good policy leads instead to a large increase in cooperation. When the good policy is selected via the random-dictator rule, the estimated increase relative to continued play under the bad policy ranges from about 73–98 percentage points in the Tax treatment, 56–82 percentage points in the PriorityTax treatment, and 86–100 percentage points in the Priority treatment. Under exogenous assignment (coin flip), the corresponding increases range from about 53–75 percentage points in the Tax treatment and 77–98 percentage points in the PriorityTax treatment. Across treatments, the increase emerges

[17]We only detect one statistically significant pre-treatment coefficient in PriorityTax ($p = 0.044$). All remaining pre-treatment coefficients are statistically insignificant ($p \geq 0.176$).

Figure 2: Event-study estimates of cooperation rates

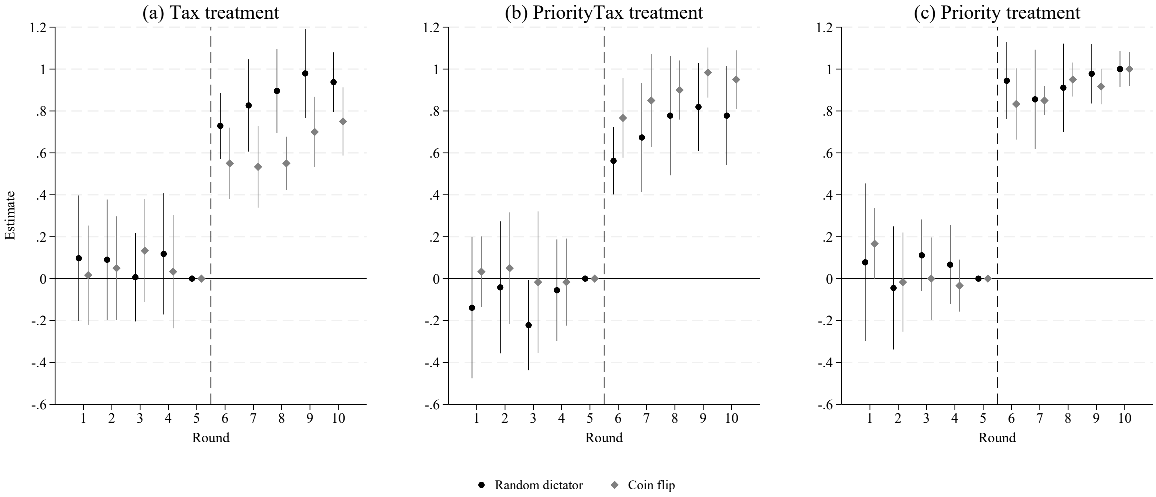

Notes: Each panel reports event-study estimates of cooperation choices. Data from Parts 1 and 2 only. The underlying specification interacts round dummies with an indicator equal to one when the game played in Part 2 is the bad policy and zero when participants instead play under the good policy. Round 5—the final round of Part 1—is the omitted category. Circles denote estimates from the random-dictator conditions; diamonds denote estimates from the corresponding coin-flip control conditions. Estimates come from linear probability models. Standard errors are clustered at the matching-group level.

immediately after the policy change and persists throughout Part 2.[18]

Results are qualitatively similar when examining efficiency, measured as the sum of pair payoffs, rather than cooperation rates (see Appendix Figure D.2). These dynamics are also visible in the raw data, as shown in Appendix Figures D.3 and D.4. Taken together, the evidence confirms that all good policies substantially increase cooperation relative to the bad policy, validating the premise of the theoretical framework in Section 3.2.

5.2 Priority policies substantially reduce support for bad policy

We now turn to whether priority policies reduce support for the bad policy. Figure 3 reports the share of participants supporting the bad policy in Part 2 by treatment, and Table 4 reports treatment-effect estimates from linear probability models. We begin by comparing the Tax and the PriorityTax treatments, where the good policies share the same tax structure. Support for the bad policy falls from 72.5% under Tax to 58.3% under PriorityTax (Figure 3).[19] The baseline specification in Column 1 of Table 4 shows that

[18]Differences in Part-2 cooperation rates between PriorityTax and Priority treatments are not statistically significant according to linear probability regressions pooling all Part 2 rounds in which the good policy is implemented (random dictator: $p = 0.951$; coin flip: $p = 0.328$). In a setting where cooperation under the priority policies is already near ceiling, cooperation under the good policy is lower in the Tax treatment, and these gaps are statistically significant (random dictator and coin flip: $p < 0.01$).

[19]In the Random Dictator treatment of Dal Bó et al. (2018), which corresponds to our Tax treatment, support for the bad policy was 52.98%. The higher level of support we observe may reflect several differences across studies. While we implemented small changes in the instructions and experimental

Figure 3: Support for bad policy in Part 2

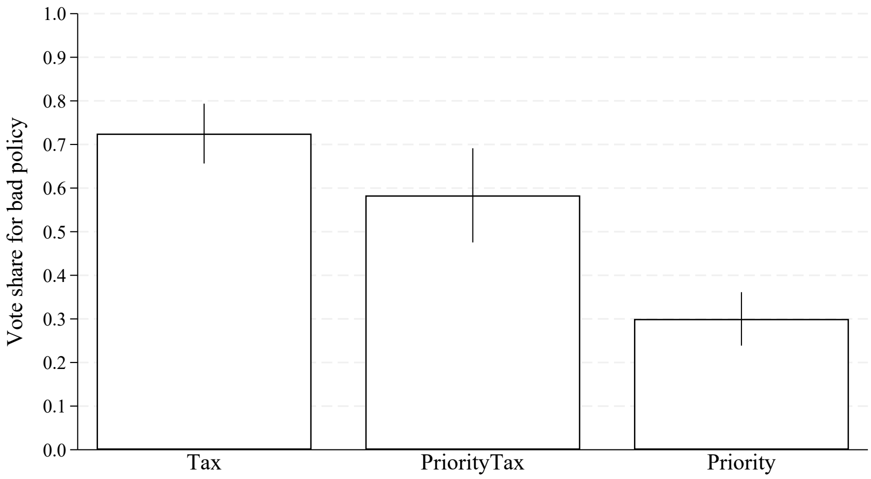

Notes: The figure reports the estimated share of participants voting for the bad policy in Part 2 across treatments. Estimates come from a linear probability model regressing the vote on treatment indicators, with standard errors clustered at the Part-1 matching group.

PriorityTax reduces support for the bad policy by 14.2 percentage points ($p = 0.031$), consistent with Prediction 1. The effect is robust to controlling for individual characteristics ($p = 0.033$ in Column 2) and to excluding participants with a poor understanding of the instructions ($p \leq 0.028$ in Columns 3 and 4).[20] The results are highly similar using probit specifications (see Appendix Table D.3).[21]

Under the Priority treatment, in line with Prediction 2, support for the bad policy falls from 58.3% to 30% relative to PriorityTax, roughly halving it. The regression analysis confirms a large and statistically significant difference between Priority and PriorityTax ($p < 0.001$).[22] In the following sections, we examine the mechanisms behind the decrease in support for the bad policy induced by the two priority treatments.

interface, one plausible explanation is subject-pool composition: our participants are mostly students at a large public university in Spain, whereas Dal Bó et al. (2018) recruited at two highly selective U.S. universities (Brown University and the University of California, Berkeley).

[20]The treatment effect remains robust, though less precisely estimated, when we additionally control for guessing-game choice and political orientation (see Appendix Table D.2). We exclude these variables from our main specifications because, following Dal Bó et al. (2018), they were elicited after treatment assignment. Moreover, some participants declined to report their political orientation, so including this control reduces the sample size.

[21]Pre-existing differences in cooperation rates under the bad policy in Part 1 are small and statistically insignificant between Tax and PriorityTax, both when pooling all five rounds and when focusing on round 5 only ($p \geq 0.752$; see Appendix Table D.1). Cooperation in Part 1 is somewhat lower under Priority than under Tax (-7.8 percentage points pooling all rounds, $p = 0.067$), but this difference, if anything, works against our main finding: lower cooperation experience should make Priority participants more pessimistic and therefore more inclined to support the bad policy.

[22]This comparison should be interpreted as the effect of the pure-priority design package: moving from PriorityTax to Priority removes taxes and increases the priority prize and punishment.

Table 4: Treatment effects: Support for bad policy in Part 2

	Full sample		Below-median mistakes	
Dep.var.: Vote for bad policy	(1)	(2)	(3)	(4)
PriorityTax	-0.142**	-0.135**	-0.179**	-0.180**
	(0.064)	(0.062)	(0.079)	(0.076)
Priority	-0.425***	-0.436***	-0.496***	-0.512***
	(0.046)	(0.051)	(0.061)	(0.070)
Controls		✓		✓
Mean dep. var. (Tax)	0.725	0.725	0.786	0.786
Wald p: PriorityTax = Priority	0.000	0.000	0.000	0.000
Observations	360	360	195	195
Matching groups	60	60	60	60

Notes: Data from Part 2 only. Estimates are from linear probability models. Standard errors (in parentheses) are clustered at the Part-1 matching group. Controls include the number of quiz attempts, age, gender, occupation, education level and field, and high-school grade. Columns (3)–(4) exclude participants whose number of mistakes in answering the quiz on the instructions exceeds the median (computed among non-control sessions); the number of quiz mistakes is absorbed by the sample restriction. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5.3 Priority policies attract pessimistic voters

The theoretical framework predicts that priority policies should reduce support for the bad policy especially among individuals who are pessimistic about others' cooperation under the good policy. Under Tax, only accurate types—those who correctly anticipate high cooperation under the good policy—should support it. Consistent with the premise of this argument, substantial shares of participants hold pessimistic beliefs about cooperation under the good policy across all treatments (Appendix Figure D.5).[23] Under the priority treatments, both accurate and pessimistic types should support the good policy (Prediction 1). This yields a testable implication: the link between beliefs about cooperation under the good policy and voting should be weaker under the priority treatments than under Tax. Thus, we use belief patterns to characterize the heterogeneity predicted by the model, rather than to separately identify the causal effect of beliefs on voting.

Figure 4 provides visual evidence. Each cell shows the share of participants voting for the good policy within a 10-point belief bin. Under Tax, support for the good policy rises steeply with beliefs: participants in the lowest bins almost never support it, while those in the highest bins mostly do. Under PriorityTax, the gradient is flatter. At nearly every belief level, support for the good policy is higher under PriorityTax than under Tax: the priority mechanism shifts voting behavior even among participants who hold the same

[23]Specifically, 52.5% of participants in Tax, 84.2% in PriorityTax, and 76.7% in Priority report beliefs below the actual cooperation rate observed in the corresponding Ctrl condition. Average beliefs about cooperation under the good policy are 68.3 in Tax, 76.5 in PriorityTax, and 78.8 in Priority; beliefs are significantly higher in both priority treatments than in Tax ($p \leq 0.005$) and similar across PriorityTax and Priority ($p = 0.329$; Appendix Table D.6).

Figure 4: Good-policy vote share by beliefs and treatment in Part 2

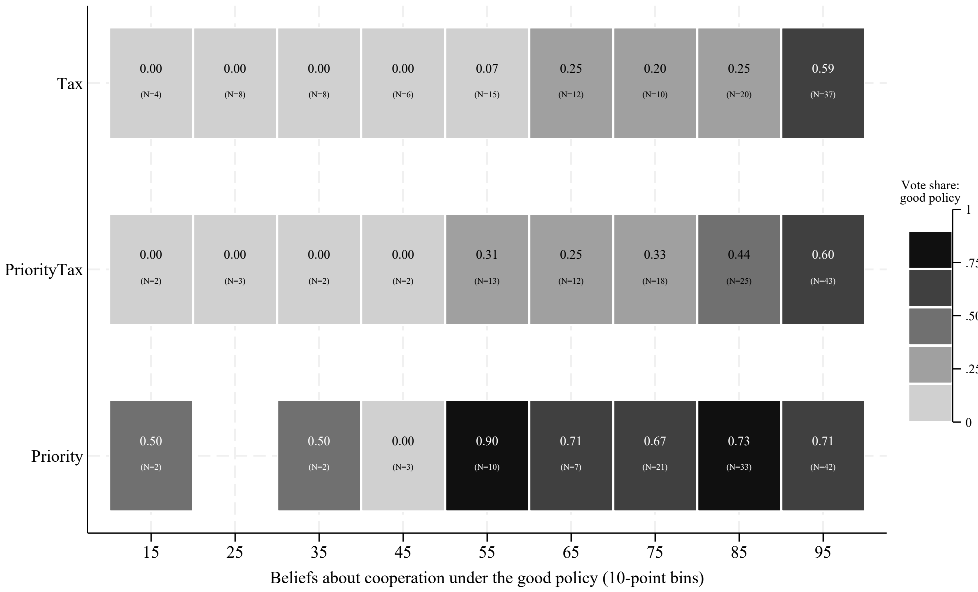

Notes: Each cell reports, for a given treatment and 10-point belief bin, the share of participants voting for the good policy; the number in parentheses gives the corresponding cell sample size. Beliefs are grouped into 10-point intervals, with bin midpoints shown on the horizontal axis. Darker shading indicates a higher vote share for the good policy. Blank cells denote treatment–belief bins with no observations.

beliefs about others' cooperation. The flatter gradient under PriorityTax is consistent with the idea that pessimists no longer need to anticipate others' cooperation to find the good policy attractive, so beliefs matter less for voting.[24] In addition, among supporters of the bad policy, most had pessimistic beliefs (65.5% in Tax, 87.1% in PriorityTax, and 77.8% in Priority).

Table 5 confirms this. In the Tax treatment, a one-percentage-point increase in expected cooperation under the good policy is associated with roughly one-percentage-point reduction in the probability of voting for the bad policy, conditional on beliefs about the bad policy (Columns 1 and 2). This gradient is reduced by about a third under PriorityTax (Columns 3 and 4). The difference in gradients between Tax and PriorityTax is statistically significant without controls ($p = 0.007$, Appendix Table D.4) and close to significance with controls ($p = 0.131$). Coefficients are essentially unchanged by the inclusion of controls.[25]

Moreover, the gradient becomes weaker still in Priority relative to PriorityTax. As shown by Corollary 1, removing taxes lowers the belief threshold above which participants

[24] Figure 4 also suggests that differences in belief levels (see Footnote 23) cannot account for the treatment effects on voting, since the priority mechanism shifts voting behavior even among participants who hold the same beliefs.

[25] Conclusions are similar if we use the belief differential $\beta^G - \beta^B$ as a single regressor: the gradient attenuates across treatments, from -0.008 ($p < 0.001$) under Tax to -0.006 ($p < 0.001$) under PriorityTax and -0.003 ($p = 0.033$) under Priority (see Appendix Table D.5). In particular, the belief-vote correlation is weaker under both priority treatments than under Tax, consistent with the pattern documented using separate beliefs.

Table 5: Beliefs and support for the bad policy in Part 2

	Tax		PriorityTax		Priority	
	(1)	(2)	(3)	(4)	(5)	(6)
Beliefs: good policy	-0.009***	-0.009***	-0.004***	-0.006***	-0.001	-0.001
	(0.001)	(0.001)	(0.001)	(0.002)	(0.003)	(0.003)
Beliefs: bad policy	0.005***	0.005***	0.009***	0.006**	0.005*	0.006*
	(0.001)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)
Controls		✓		✓		✓
Belief good − bad	-0.014	-0.015	-0.013	-0.012	-0.006	-0.007
	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)
Observations	120	120	120	120	120	120
Clusters	20	20	20	20	20	20
Adjusted R^2	0.278	0.258	0.247	0.276	0.031	-0.007

Notes: Data from Part 2 only. Estimates are from linear probability models. Standard errors (in parentheses) are clustered at the Part-1 matching group. Belief variables are measured on a 0–100 scale, with 0 indicating no expected cooperation and 100 indicating certainty about cooperation. Controls include the number of quiz attempts, age, gender, occupation, education level and field, and high-school grade. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

support the good policy, so the additional support gained by moving from PriorityTax to Priority should come disproportionately from participants with low beliefs about others' cooperation (Prediction 3). Consistently, Figure 4 shows that the increase in support for the good policy when moving from PriorityTax to Priority is concentrated in the lower belief bins, whereas support among optimistic participants is already high under both treatments. As a result, the gradient under Priority is essentially zero and statistically insignificant (Table 5, Columns 5 and 6). A premise of Corollary 1 is that beliefs about cooperation under the bad policy are not accurate. This is indeed the case, as shown later in Section 5.5.

Table 5 also reveals that beliefs about cooperation under the bad policy are positively associated with voting for the bad policy in the Tax, PriorityTax, and Priority treatments: participants who are more optimistic about cooperation under the status quo are more inclined to retain it. We return to this channel in Section 5.5, after first examining the role of inertial participants in Section 5.4.

5.4 A limited role for inertia

While PriorityTax reduces support for the bad policy relative to Tax, a majority still prefers the bad policy. Consistent with Prediction 4, a possible explanation for the limited effectiveness of PriorityTax is the presence of a substantial share of inertial participants, i.e., those who fail to anticipate that their own behavior will change under the new policy. These participants underestimate both others' and their own willingness to cooperate under the good policy and consequently view the bad policy as more attractive: in equilibrium it yields a payoff of 5, whereas they expect to get 1 under the good policy.

We did not elicit beliefs about participants' own cooperation, given the complexity of doing so, and the risk of influencing later individual cooperation. Instead, we implemented the PriorityTaxPlus treatment, which is identical to PriorityTax except that the mutual-defection payoff under the good policy is increased to 5. This modification makes inertial participants indifferent between the bad and the good policy. Moreover, by equalizing mutual-defection payoffs across policies, it may lead participants to place greater weight on other outcomes, potentially increasing support for the good policy.

Although the modified priority policy generates sizable gains in cooperation and efficiency (see Appendix Figure D.6), the drop in support for the bad policy is modest: 50.8% of participants vote for the bad policy under PriorityTaxPlus, compared to 58.3% under PriorityTax. The regression estimate is less than 8 percentage points and not statistically significant ($p \geq 0.296$; Columns 1–2 of Appendix Table D.7). Among participants with a good understanding of the instructions, the effect rises to about 18 percentage points ($p \leq 0.043$; Columns 3–4).

These estimates provide a lower bound on the share of participants whose support for the bad policy is driven by inertial reasoning.[26] The implied share of inertial participants—between 8 and 18 percentage points of participants, depending on the sample—is limited relative to the overall demand for the bad policy under PriorityTax (58.3%). Inertial types thus contribute to PriorityTax's limited effectiveness but are not its primary obstacle, consistent with the conclusion of Dal Bó et al. (2018) that the demand for bad policy is primarily driven by the failure to anticipate others' behavioral adjustment.

5.5 Overoptimism about the status quo predicts support for the bad policy

In Section 5.3, we discussed how beliefs about cooperation under the good policy correlate with support for the bad policy, holding beliefs about the bad policy constant. Here we focus instead on beliefs under the bad policy (Table 5). Our baseline theoretical analysis assumes that voters hold equilibrium beliefs under the bad policy, i.e., that they expect no cooperation (Assumption 1). However, inequality (2) shows that if voters are optimistic about cooperation under the bad policy ($\beta^B > 0$), the belief threshold about cooperation under the good policy above which they support reform increases: a higher perceived payoff under the status quo makes reform less attractive.

Despite five rounds of experience with the bad policy, a large share of participants remains surprisingly optimistic about cooperation under it. Figure 5 illustrates that

[26] The bound is conservative because PriorityTaxPlus makes inertials indifferent between policies rather than strictly preferring the good one, so some inertials may still vote for the bad policy.

Figure 5: Beliefs under bad policy in Part 2

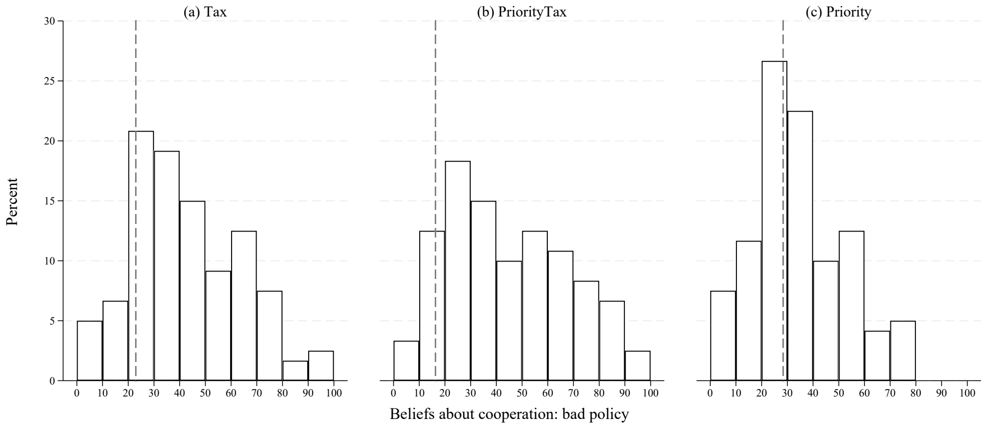

Notes: Each panel shows the distribution of beliefs about cooperation under the bad policy in Part 2 by treatment. The histograms report the percentage of participants in 10-point bins of stated beliefs (0–100). The vertical segment marks the actual cooperation rate under the bad policy in Part 2 for the corresponding control condition.

beliefs about cooperation under the bad policy frequently exceed the cooperation rates observed in the corresponding Ctrl conditions.[27] Appendix Figure D.7 suggests that this overoptimism partly reflects the projection of one's own past behavior onto others: participants who cooperated more frequently in Part 1 hold more optimistic beliefs about cooperation under the bad policy.

This overoptimism is positively associated with support for the bad policy. In all treatments, the large majority of bad-policy supporters are overoptimistic (79.3% in Tax, 92.9% in PriorityTax, and 69.4% in Priority). Moreover, in all three treatments, a ten-percentage-point increase in expected cooperation under the bad policy is associated with a 5 to 9 percentage point increase in the probability of voting for the bad policy ($p \leq 0.075$; Table 5), with only modest differences across treatments ($p \geq 0.069$ without controls, $p \geq 0.332$ with controls; Appendix Table D.4).

These patterns suggest that overoptimism about cooperation under the bad policy is an additional driver of demand for the bad policy, distinct from the pessimism about the good policy documented in Section 5.3. Pooling across treatments, nearly two-thirds of bad-policy voters are simultaneously pessimistic about the good policy and overoptimistic about the bad policy, suggesting that the two biases reinforce each other in sustaining demand for the bad policy. Unlike pessimism about the good policy, overoptimism about the bad policy is not directly targeted by the priority mechanism, which helps explain why residual demand for the bad policy persists even under the Priority treatment.

[27]Specifically, 78.3% of participants in Tax, 85.8% in PriorityTax, and 54.2% in Priority report beliefs that exceed the actual cooperation rate. Relative to the corresponding Ctrl benchmarks, participants overestimate cooperation under the bad policy by 17.4 percentage points in Tax and 25.8 percentage points in PriorityTax. The difference between these two overestimation gaps is not statistically significant ($p = 0.283$). Overoptimism is less pronounced under Priority, where beliefs are significantly lower ($p \leq 0.006$; Appendix Table D.6).

The intuition in Dal Bó et al. (2018) is that players carry over their beliefs from a game with which they have experience to a new game. We can formalize this intuition with a theory that rationalizes ex-post the main experimental findings, namely pessimism under the good policy, optimism under the bad policy, and the monotonic decline in support for the bad policy from Tax to PriorityTax to Priority. Following the spirit of Bordalo et al. (2020), our theory rests on two assumptions about beliefs. First, a player's belief about the counterpart's cooperation is a weighted average of (i) the counterpart's best response and (ii) the counterpart's past actions. Second, a player's belief about her own cooperation is a weighted average of (i) her own best response and (ii) her own past actions. Initial beliefs about both self and the counterpart are set to 1/2, consistent with the initial cooperation rates observed in the experiment under the bad policy. The key feature of the theory is that, when a player faces a new game (such as any of our good policies), past actions, both her own and the counterpart's, feed into her beliefs about cooperation in the new game. Formally, beliefs are defined as follows:

$$\alpha_t^H \;=\; k_a \cdot \frac{\sum_{z=0}^{t-1} \delta^{t-1-z}\, a^z}{\sum_{z=0}^{t-1} \delta^{t-1-z}} \;+\; (1-k_a) \cdot \alpha_{BR}^H, \tag{3}$$

$$\beta_t^H \;=\; k_b \cdot \frac{\sum_{z=0}^{t-1} \delta^{t-1-z}\, b^z}{\sum_{z=0}^{t-1} \delta^{t-1-z}} \;+\; (1-k_b) \cdot \beta_{BR}^H, \tag{4}$$

where $H \in \{B, G\}$ denotes the game under consideration; $a^z \in \{0,1\}$ and $b^z \in \{0,1\}$ are own and counterpart's actions in period z, with 1 denoting cooperation; and α_{BR}^H and β_{BR}^H are the best-response beliefs for self and other in game H. Even though actions can only take values in $\{0,1\}$, with abuse of notation we assume $a^0 = b^0 = 1/2$, which implies that players with $k_a = k_b = 1$ start with the anchor that they and the counterpart would randomize equally across actions. This is consistent with the initial observed cooperation rates in the experiment and evokes the concept of level-0 thinking. The parameter k_a captures inertia with respect to a player's own past plays, while k_b captures the analogous inertia about others. Finally, $\delta \in [0,1]$ discounts past actions.

A player with $k_a = k_b = 0$ is an accurate player. To see how the theory generates pessimism, consider a PD played for n rounds between players A and B. If player B has $k_a < 1$, she will eventually start defecting. This defection feeds into player A's belief β^{PD}, which then drifts downward over time. Now suppose player A faces an alternative new game (say, a good policy). If she has low k_a and high k_b, she is empirically a pessimist: her belief about the counterpart's cooperation in the new game largely reflects the counterpart's past defections in the PD.

Simulations reported in Appendix Figure D.1 with continuous distributions of k_a and k_b across the population yield two key results. First, the distribution of β^G at the voting stage is not concentrated at 1, i.e., pessimists exist. Second, the distribution of β^{PD} at the voting stage is not concentrated at 0 i.e., players exhibit overoptimism under the PD.

Moreover, simulations in table B.1 in the Appendix with different continuous distributions of k_a and k_b on $[0, 1]$ show that the support for the PD decreases as we move from the Tax treatment to the PriorityTax treatment and finally to the Priority treatment.

This theory therefore provides a behavioral microfoundation for the simplified theory presented above and accommodates the main empirical regularities observed in the experiment, including the overoptimism under the bad policy.

5.6 Status-quo bias cannot explain the treatment effects

One factor not incorporated into our theoretical framework that could contribute to support for the bad policy is status-quo bias, whereby participants may prefer to retain the policy under which they are currently playing rather than switch to an alternative. Two results suggest that this factor appears to play at most a secondary role. First, since our treatments are identical in all design features except the payoff structure of the good policy, a generic preference for the status quo cannot explain the differential treatment effects reported in Section 5.2.

Second, support for the bad policy remains substantial in the RevPriorityTax treatment, in which participants initially play under the good policy (with the same payoffs as PriorityTax) and subsequently vote on whether to retain it or switch to the bad policy. In this treatment, 60.8% vote for the bad policy, compared to 58.3% under PriorityTax; the difference is not statistically significant ($p \geq 0.403$; see Appendix Table D.7). This is true even though switching to the bad policy leads to large and statistically significant declines in cooperation and efficiency (see Appendix Figure D.8).[28]

5.7 Learning to support the good policy

In all treatments except Priority a majority of the participants votes for the bad policy in Part 2. We next examine whether learning mitigates this pattern by comparing voting behavior between Parts 2 and 3. Figure 6 reports the estimated change in support for the bad policy by treatment. In the full sample (circles), support for the bad policy declines significantly in Part 3 under Tax and PriorityTax—by between 13 and 18 percentage points ($p \leq 0.001$)—while it remains statistically unchanged under Priority ($p = 0.627$). When focusing on participants who experienced the good policy in Part 2 (diamonds), the decline in support for the bad policy becomes larger in all treatments except Tax, where

[28]Predicting the treatment effect for RevPriorityTax is difficult, because the presence of pessimist and inertial participants, and of taxes t_D and t_C would increase the support for the bad policy, while the presence of accurate participants, and a status quo bias, would increase the support for the good policy. Assuming no status quo bias, there should be a total increase in support for the bad policy, with respect to PriorityTax. A small status quo bias would make the two treatments similar in terms of support for the bad policy.

Figure 6: Support for the bad policy: Part 2 vs. Part 3

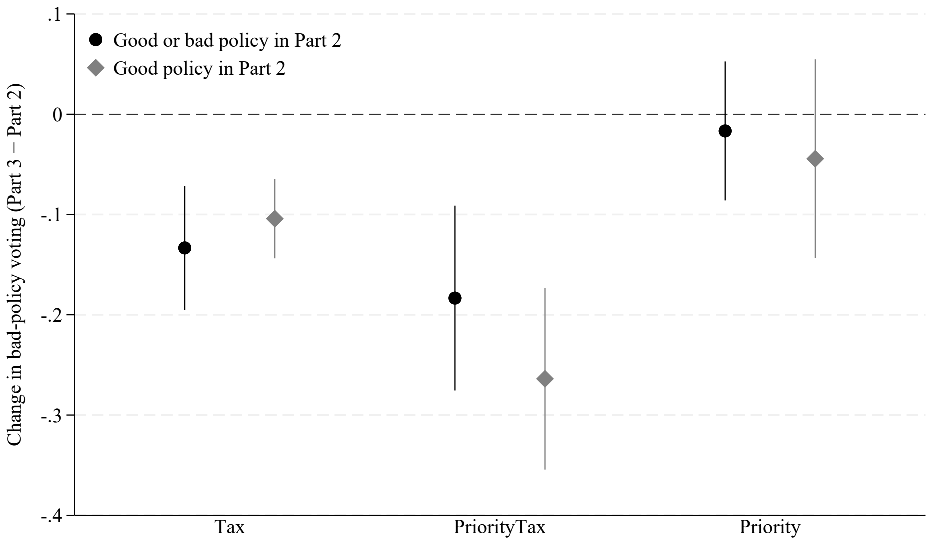

Notes: The figure reports estimates from linear probability models of voting for the bad policy. The dependent variable equals one if the participant votes for the bad policy. Specifications include treatment indicators interacted with a Part 3 indicator and controls for the number of quiz attempts, age, gender, occupation, education level and field, and high-school grade. Reported coefficients correspond to within-treatment changes in voting behavior between Part 2 and Part 3. Circles denote estimates from the full sample; diamonds restrict to participants who experienced the good policy in Part 2. Standard errors are clustered at the Part-2 matching-group level.

it becomes slightly less pronounced.[29] Finally, restricting the sample to participants who voted for the bad policy in Part 2, experiencing the good policy is associated with a reduction in the probability of voting for the bad policy again of 9 percentage points under Tax ($p = 0.186$), 21 under PriorityTax ($p = 0.194$), and 51 under Priority ($p = 0.021$), as reported in Appendix Table D.8.

Patterns in belief data further highlight the importance of experience. To study belief updating from Part 2 to Part 3, we pool data from Tax, PriorityTax, and Priority. We focus on participants who voted for the bad policy. The first panel of Figure 7a shows that, among these participants, those who experienced the good policy in Part 2 revise their beliefs about cooperation under the good policy upward by about 15 percentage points ($p < 0.001$; see Appendix Table D.9). The second panel of Figure 7a shows that participants who again experienced the bad policy in Part 2 revise their bad-policy beliefs downward by about 10 percentage points ($p < 0.001$). In contrast, beliefs about the policy participants did not experience in Part 2 exhibit little change in either case ($p \geq 0.674$). Thus, experience generates belief updating in the direction implied by realized outcomes, but does not spill over to beliefs about the alternative policy. Notably, beliefs about cooperation under the bad policy in Part 3 remain elevated—around 35–40%—even after updating.

[29] By design, experiencing the good policy in Part 2 is at least partially endogenous. Therefore, causal interpretation warrants some caution.

Figure 7: Belief updating

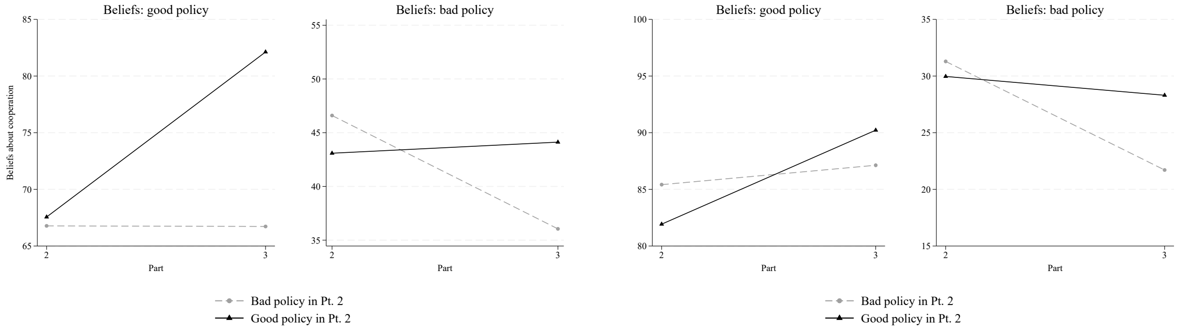

(a) Voted bad policy in Part 2 (b) Voted good policy in Part 2

Notes: The figure reports average beliefs about cooperation in Parts 2 and 3, separately by the policy played in Part 2. Data are from the Tax, PriorityTax, and Priority treatments. The sample is restricted to participants who voted for the bad policy in Part 2 (left panel) and for the good policy in Part 2 (right panel). Beliefs are measured on a 0–100 scale. Each point represents the mean belief within a Part × policy-played cell.

Turning to participants who voted for the good policy in Part 2, Figure 7b shows that their beliefs are more accurate overall and, consequently, updating is less pronounced. The largest change concerns beliefs under the bad policy among those who again experienced the bad policy in Part 2: these beliefs decline by about 10 percentage points, reaching approximately 20% ($p = 0.001$). Beliefs about cooperation under the good policy are already above 80% for those who experienced the good policy in Part 2 and rise to around 90% in Part 3 ($p < 0.001$). As in the previous case, beliefs about the policy not experienced in Part 2 again display only minor variation ($p \geq 0.314$).

In general, learning from experience helps reduce demand for the bad policy under Tax and PriorityTax, but not under Priority, where support for the bad policy is already at its minimum across treatments—albeit still around 30%—suggesting a floor on the demand for bad policy that even well-designed remedies cannot fully eliminate.

6 Conclusion

We conclude that part of the electorate's demand for bad policy is a policy-design problem. When voters underweight equilibrium responses, the ballot box can favor arrangements that feel privately safe yet are socially inefficient. Our remedy is to redesign the reform so that cooperation is privately attractive even for pessimistic voters: those who cooperate when others defect receive a priority benefit, while those who defect when others cooperate are penalized. In the laboratory, replacing the bad-policy status quo with any good policy immediately and persistently raises cooperation and joint earnings. But political support differs sharply across designs. Adding priority incentives to a standard tax reduces votes for the bad policy by about fourteen percentage points;

the pure-priority, which removes taxes and implements larger budget-balanced priority incentives, reduces support for the bad policy to thirty percent.

These findings carry clear implications for policy design. First, policies should be tailored to pessimistic voters: the priority prize and punishment should be set high enough to offset the private cost of cooperation and, when voters are overoptimistic about the status quo, the perceived opportunity cost of abandoning it. Second, Pigouvian taxes should be used sparingly. When optimism about the status quo is salient, they can make reforms look less attractive ex ante, whereas prize-punishment priority schemes preserve efficiency while securing greater political acceptance. Third, policymakers should leverage learning: pilot programs and phased rollouts that make others' behavioral adjustments salient can durably shift beliefs toward efficient expectations and, with them, electoral support.

Many contested reforms, including environmental regulation, require citizens to anticipate how others will adjust their behavior. Carbon taxes, for example, are efficient but often face strong public opposition (Umit and Schaffer, 2020), which may partly reflect voters' difficulty in anticipating their behavioral and social effects. Priority-like features offer a different route to political feasibility. Degressive subsidies for green technologies, such as early-adopter rebates for solar panels, make participation especially attractive when few others have adopted; as adoption grows, the private priority advantage fades. This degressive structure insures early movers against the risk that others will not follow—precisely the property that drives political support in our experiment. Whether such instruments can be deliberately adapted to large-scale regulation remains an open question. The design principle, however, is general: reforms are more likely to survive political selection when they do not require voters to trust that others will adjust.

References

Acemoglu, D., Ozdaglar, A., and Siderius, J. (2024). A model of online misinformation. *Review of Economic Studies*, 91(6):3117–3150.

Algan, Y. and Cahuc, P. (2010). Inherited trust and growth. *American Economic Review*, 100(5):2060–2092.

Ali, S. N., Mihm, M., and Siga, L. (2025). The political economy of zero-sum thinking. *Econometrica*, 93(1):1–36.

Angelucci, C. and Prat, A. (2024). Is journalistic truth dead? Measuring how informed voters are about political news. *American Economic Review*, 114(4):887–925.

Arteaga, F., Kapor, A. J., Neilson, C. A., and Zimmerman, S. D. (2022). Smart matching platforms and heterogeneous beliefs in centralized school choice. *Quarterly Journal of Economics*, 137(3):1791–1848.

Ashkenazi, T., Lavee, J., and Mor, E. (2015). Organ donation in Israel—Achievements and challenges. *Transplantation*, 99(2):265–266.

Besley, T. (2005). Political selection. *Journal of Economic Perspectives*, 19(3):43–60.

Bigoni, M., Bortolotti, S., Casari, M., and Gambetta, D. (2019). At the root of the North–South cooperation gap in Italy: Preferences or beliefs? *Economic Journal*, 129(619):1139–1152.

Blanes i Vidal, J., Draca, M., and Fons-Rosen, C. (2012). Revolving door lobbyists. *American Economic Review*, 102(7):3731–3748.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2020). Memory, attention, and choice. *Quarterly Journal of Economics*, 135(3):1399–1442.

Bursztyn, L., González, A. L., and Yanagizawa-Drott, D. (2020). Misperceived social norms: Women working outside the home in Saudi Arabia. *American Economic Review*, 110(10):2997–3029.

Bursztyn, L. and Yang, D. Y. (2022). Misperceptions about others. *Annual Review of Economics*, 14(1):425–452.

Caselli, F. and Morelli, M. (2004). Bad politicians. *Journal of Public Economics*, 88(3-4):759–782.

Chandler, J. A., Burkell, J. A., and Shemie, S. D. (2012). Priority in organ allocation to previously registered donors: Public perceptions of the fairness and effectiveness of priority systems. *Progress in Transplantation*, 22(4):413–422.

Charness, G. and Rabin, M. (2002). Understanding social preferences with simple tests. *Quarterly Journal of Economics*, 117(3):817–869.

Chen, D. L., Schonger, M., and Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97.

Dal Bó, E., Dal Bó, P., and Eyster, E. (2018). The demand for bad policy when voters underappreciate equilibrium effects. *Review of Economic Studies*, 85(2):964–998.

Dal Bó, P., Foster, A., and Putterman, L. (2010). Institutions and behavior: Experimental evidence on the effects of democracy. *American Economic Review*, 100(5):2205–2229.

Diamond, D. W. and Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, 91(3):401–419.

Douenne, T. and Fabre, A. (2022). Yellow Vests, pessimistic beliefs, and carbon tax aversion. *American Economic Journal: Economic Policy*, 14(1):81–110.

Dreyfuss, B. (2026). Equilibrium neglect and political feasibility. Mimeo.

Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, 114(3):817–868.

Fernandez, R. and Rodrik, D. (1991). Resistance to reform: Status quo bias in the presence of individual-specific uncertainty. *American Economic Review*, 81(5):1146–1155.

Folke, O. (2014). Shades of brown and green: Party effects in proportional election systems. *Journal of the European Economic Association*, 12(5):1361–1395.

Gagliarducci, S. and Nannicini, T. (2013). Do better paid politicians perform better? Disentangling incentives from selection. *Journal of the European Economic Association*, 11(2):369–398.

Ghidoni, R., Cleave, B. L., and Suetens, S. (2019). Perfect and imperfect strangers in social dilemmas. *European Economic Review*, 116:148–159.

Glaeser, E. L. and Ponzetto, G. A. (2017). Fundamental errors in the voting booth. NBER Working Paper No. 23683.

Hughes, J. E. and Podolefsky, M. (2015). Getting green with solar subsidies: Evidence from the California Solar Initiative. *Journal of the Association of Environmental and Resource Economists*, 2(2):235–275.

Karni, E. (2009). A mechanism for eliciting probabilities. *Econometrica*, 77(2):603–606.

Kartal, M. and Tyran, J.-R. (2022). Fake news, voter overconfidence, and the quality of democratic choice. *American Economic Review*, 112(10):3367–3397.

Kessler, J. B. and Roth, A. E. (2012). Organ allocation policy and the decision to donate. *American Economic Review*, 102(5):2018–2047.

Kim, J., Li, M., and Xu, M. (2021). Organ donation with vouchers. *Journal of Economic Theory*, 191:105159.

Knack, S. and Keefer, P. (1997). Does social capital have an economic payoff? A cross-country investigation. *Quarterly Journal of Economics*, 112(4):1251–1288.

Kroft, K., Laliberté, J.-W., Leal-Vizcaíno, R., and Notowidigdo, M. J. (2024). Salience and taxation with imperfect competition. *Review of Economic Studies*, 91(1):403–437.

Lee, D. S., Moretti, E., and Butler, M. J. (2004). Do voters affect or elect policies? Evidence from the U.S. House. *Quarterly Journal of Economics*, 119(3):807–859.

Li, S. (2017). Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–3287.

Luechinger, S. and Moser, C. (2014). The value of the revolving door: Political appointees and the stock market. *Journal of Public Economics*, 119:93–107.

Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5):1313–1326.

North, D. C. (1990). *Institutions, Institutional Change and Economic Performance*. Cambridge University Press.

Nunnari, S., Proto, E., and Rustichini, A. (2024). Cognitive abilities and the demand for bad policy. CESifo Working Paper No. 11206.

Roberti, P. (2019). Citizens or lobbies: Who controls policy? *Games and Economic Behavior*, 113:497–514.

Robinson, J. A. and Acemoglu, D. (2012). *Why nations fail: The origins of power, prosperity and poverty*. Profile London.

Sutter, M., Haigner, S., and Kocher, M. G. (2010). Choosing the carrot or the stick? Endogenous institutional choice in social dilemma situations. *Review of Economic Studies*, 77(4):1540–1566.

Umit, R. and Schaffer, L. M. (2020). Attitudes towards carbon taxes across Europe: The role of perceived uncertainty and self-interest. *Energy Policy*, 140:111385.

Online appendix

A Extension to N players

In this section we study the robustness of our theory to two realistic modifications of our baseline model. We consider, first, the case of n players and second, non-linear production functions for the public good. To make the n-player version of the Prisoner's Dilemma comparable with the two-player version of the main theory, we assume that the defectors take all the benefits from cooperation. However, the results are unaffected by considering an n-player version of the Prisoner's Dilemma where the pool of resources is distributed equally to all players, independently from their cooperation decision.

N-player Prisoner's Dilemma. We consider an n-player Prisoner's Dilemma, where each cooperator adds a benefit b to a pool of resources. This pool of resources is distributed equally only among the defectors. Cooperation entails a cost c. Thus, the payoff from cooperation, with n_C cooperators and at least one defector, is $-c$, while the payoff from defection is $n_C b/(n - n_C)$. When all players cooperate, the payoff each of them gets is $b - c$. It is immediate to see that defection is a dominant strategy. As in the main theory, we assume $b \geq c$, cooperation maximizes total surplus.

N-player Tax Game. Adapting the Tax game to n players is straightforward. The payoff from cooperation, with at least one defector, is $-t_C - c$, while the payoff from defection is $n_C b/(n - n_C) - t_D$. The payoff when everyone cooperates is $b - c - t_C$. In this case full cooperation is an equilibrium if $b - c - t_C \geq (n-1)b - t_D$.

N-player Priority Policy Game. Finally, we extend the PPG to n players, including a prize and a punishment that depend on the number of cooperators and defectors. In order to endogenize prizes and punishments and to guarantee budget balance for the policies at hand, we assume that the prize P, given to a single cooperator—where a total of n^c players cooperate—is a share $\gamma \leq 1$ taken from the $n - n_C$ defectors, of the benefit b added to society by her cooperation: γb. Similarly, the punishment p is a share $\delta \leq 1$ of the benefit b extracted by the defector from the n_C cooperators: $-\delta b n_C/(n - n_C)$. It is immediate to see that budget balance is guaranteed if $-n_C \gamma b + (n - n_C)\delta b n_C/(n - n_C) \geq 0$ or $\gamma \leq \delta$. Consequently, the payoff of a cooperator in the priority policy is $\gamma b - c$, which is equal to $b - c$ if $\gamma = 1$. The payoff of a cooperator is $b - c$ also if everyone cooperates. The payoff of a defector is $(1-\delta) b n_C/(n - n_C)$, equal to 0 if $\delta = 1$. Full cooperation is an equilibrium in the PPG if $b - c \geq (1-\delta) b (n-1)$. This inequality is satisfied if δ is sufficiently large.

Let us now characterize the game preferred by players characterized by the bias examined in Dal Bó et al. (2018) and in our main text. First, we redefine β as the expected number $\bar{n}_C$ of other players cooperating, while α is as before the own belief of cooperation.[30] Let us compare the Prisoner's Dilemma and the Tax game. A player prefers the Prisoner's Dilemma to the Tax game if

$$-\alpha^B c + (1-\alpha^B)\frac{\bar{n}_C^B}{n-\bar{n}_C^B} b \geq -\alpha^G (t_C + c) + (1-\alpha^G) b \left(\frac{\bar{n}_C^G}{n-\bar{n}_C^G} - t_D\right)$$

As in Dal Bó et al. (2018), if $\bar{n}_C^G$ is sufficiently low, the player has a preference for the Prisoner's Dilemma over the Tax game.

Let us now compare a player's preference for the PPG to the Prisoner's Dilemma. This is the case if

$$-\alpha^B c + (1-\alpha^B)\frac{\bar{n}_C^B}{n-\bar{n}_C^B} b \leq \alpha^G (\gamma b - c) + (1-\alpha^G) b (1-\delta)\frac{\bar{n}_C^G}{n-\bar{n}_C^G},$$

when $\bar{n}_C^G < n-1$. The inequality becomes

$$-\alpha^B c + (1-\alpha^B)\frac{\bar{n}_C^B}{n-\bar{n}_C^B} b \leq \alpha^G (b - c) + (1-\alpha^G) b (1-\delta)(n-1),$$

when $\bar{n}_C^G = n-1$. If the player believes that all participants will play the Nash equilibrium in both games ($\alpha^B = 0, \bar{n}_C^B = 0, \alpha^G = 1, \bar{n}_C^G = n-1$), then the inequality is satisfied. If instead the player holds belief $\bar{n}_C^G$ sufficiently low, then the inequality is satisfied if α^B and $\bar{n}_C^B$ are sufficiently low,[31] α^G is sufficiently large and γ is sufficiently large, which means that the player is a *pessimist* and the priority prize is large. In order for the priority prize to have a bite, the previous inequality must be satisfied, for some parameters and beliefs, if and only if γ is sufficiently large. This is true when, e.g., $\alpha^B = 0, \bar{n}_C^B = 0, \alpha^G = 1, \bar{n}_C^G = 0$, for which the inequality reduces to $\gamma b \geq c$, inequality which is not satisfied for $\gamma = 0$, and is satisfied for $\gamma = 1$. This proves the n-player version of Proposition 1.

A.1 Non-linear social surplus

So far, the total amount of social surplus was assumed to be a linear function of the number of cooperators: $B(n_C) = bn_C$. Arguably, in many real life examples different production functions for the public good might be more descriptive of reality. Then, to test the robustness of our results consider now a different production function for

[30] Notice that this is a degenerate distribution on the number of other players cooperating, which imposes probability 1 to $\bar{n}_C$ cooperating. An alternative would be to have a non-degenerate distribution on $\bar{n}_C$. However, this approach would make comparative statics more complex, without adding intuition to the results.

[31] As in the main text, these beliefs summarize the experience of the player playing the PD and converging to the dominant strategy equilibrium of defection by all players.

the public good, where each cooperator increases the total amount of social surplus $B(n_C)$ less than linearly, so that $B(n_C) - B(n_C - 1)$ decreases with n_C. We assume $B(0) = 0$. The rules for the Prisoner's Dilemma are the same as with linear surplus. Notice that either i) the optimization of surplus requires that there exists $n_C < n$ such that $B(n_C) - B(n_C - 1) \geq c$, and $B(n_C + 1) - B(n_C) < c$ i.e., the increase in total benefit given by n_Cth cooperator is at least equal to the cost of her cooperation, while adding an additional cooperator is inefficient; or ii) cooperation by all players is efficient: $B(n) - B(n-1) \geq c$.

Consistent with the previous section the prize P given to a cooperator—when n_c players cooperate—is a share $\gamma \leq 1$, taken from the $(n - n_C)$ defectors, of the individual benefit $B(n_C)/n_C$ for each player when she cooperates, that is $\gamma B(n_C)/n_C$. γ is set to 0 when an efficient number of cooperators cooperates.

Similarly, the punishment p is a share $\delta \leq 1$, of the benefit extracted by the defector from the n_C cooperators, that is $-\delta B(n_C)/(n - n_C)$. Again budget feasibility is guaranteed if $-n_C \gamma B(n_C)/n_C + (n - n_C)\delta B(n_C)/(n - n_C) \geq 0$, i.e., $\gamma \leq \delta$. Consequently, the payoff of a cooperator in the PPG is $\gamma B(n_C)/n_C - c$ while the payoff of a defector is $(1 - \delta)B(n_C)/(n - n_C)$. δ is set to 0 when an efficient number of cooperators cooperates.

Now, for simplicity, assume that surplus is maximized when all players cooperate: $B(n) - B(n-1) \geq c$.[32] Then, full cooperation is an equilibrium if $B(n)/n - c \geq (1 - \delta)B(n - 1)$. Let us check whether, with the largest feasible punishment, $\delta = 1$, full cooperation can be sustained. The inequality becomes $B(n) \geq nc$, which is a necessary condition for the maximization of surplus when all players cooperate.

Let us now compare a player's preference for the non-linear PPG to the Prisoner's Dilemma. This is the case if

$$-\alpha^B c + (1 - \alpha^B)\frac{B(\bar{n}_C^B)}{n - \bar{n}_C^B} \leq$$
$$\alpha^G \left[\gamma B(\bar{n}_C^G + 1)/(\bar{n}_C^G + 1) - c\right] + (1 - \alpha^G)(1 - \delta)B(\bar{n}_C^G)/(n - \bar{n}_C^G),$$

when $\bar{n}_C^G < n - 1$. This condition becomes

$$-\alpha^B c + (1 - \alpha^B)\frac{B(\bar{n}_C^B)}{n - \bar{n}_C^B} \leq$$
$$\alpha^G \left[B(\bar{n}_C^G + 1)/(\bar{n}_C^G + 1) - c\right] + (1 - \alpha^G)(1 - \delta)B(\bar{n}_C^G)/(n - \bar{n}_C^G),$$

when $\bar{n}_C^G = n-1$. If the player believes that all participants will play the Nash equilibrium in both games ($\alpha^B = 0, \bar{n}_C^B = 0, \alpha^G = 1, \bar{n}_C^G = n - 1$), then the inequality is satisfied. Now consider $\bar{n}_C^G$ sufficiently low. The previous inequality is satisfied if α^B and $\bar{n}_C^B$ are sufficiently low, α^G is sufficiently large and γ is sufficiently large, which means, as before,

[32]The same computations apply when the surplus is maximized for $n_C < n$.

that the player is a *pessimist* and the priority prize is large.

In order for the priority prize to have a bite, the previous inequality must be satisfied, for some parameters and beliefs, if and only if γ is sufficiently large. This is true when, e.g., $\alpha^B = 0, \bar{n}_C^B = 0, \alpha^G = 1, \bar{n}_C^G = 0$, for which the inequality reduces to

$$\gamma B(1) - c \geq 0,$$

and it is not satisfied for $\gamma = 0$, while it is satisfied for $\gamma = 1$. This proves the non-linear n-player version of Proposition 1.

B Simulated history-dependent beliefs

The following graph shows the simulated distributions of beliefs for the counterpart under the bad and the good policy, with k_a and k_b independently distributed as $Beta(0.5, 0.8)$.

Figure D.1: Simulated beliefs

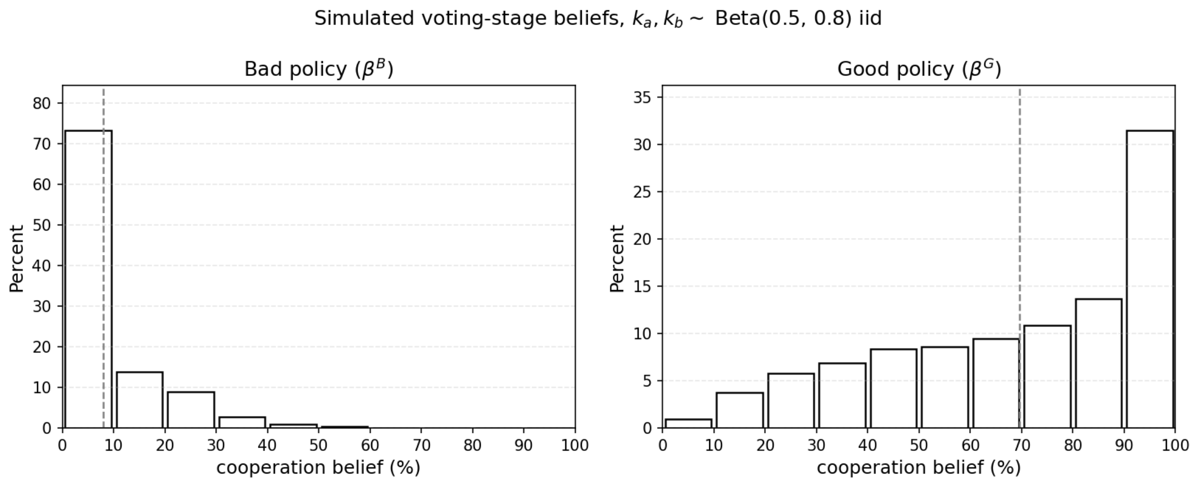

Simulations of support for the bad policy under Tax, Prioritytax, and Priority treatments, with different distributions of k_a and k_b.

The following table shows the simulations of support for the bad policy under Tax, Prioritytax, and Priority treatments, with different distributions of k_a and k_b, at the beginning of period 6. The parameter k_a captures inertia with respect to a player's own past plays, while k_b captures the analogous inertia about others. For each distribution, we draw 15,000 independent matching groups of 6 players, run the recursion through 5 Part-1 PD periods using $\alpha^0 = \beta^0 = 1/2$ and $\delta = 1$, and compute each voter's voting margin under the experimental payoff matrices. Table B.1 reports the resulting population shares voting for the bad policy across the three good-policy treatments. The first block uses i.i.d. symmetric distributions for k_a and k_b on $[0, 1]$; the second block enforces $k_a \leq k_b$ pointwise via a hierarchical truncated-normal design in which $k_a \mid k_b$ is truncated to $[0, k_b]$

with mean $k_b/2$ and standard deviation $k_b/4$. The predicted monotonicity holds in every calibration.

Table B.1: Simulated PD shares across simulations of (k_a, k_b). All calibrations satisfy the predicted monotonicity $\Pr(\text{PD} \mid \text{Tax}) \geq \Pr(\text{PD} \mid \text{PTax}) \geq \Pr(\text{PD} \mid \text{Pri})$.

Distribution of (k_a, k_b)	Pr(PD|Tax)	Pr(PD|PTax)	Pr(PD|Pri)
i.i.d. symmetric distributions on $[0, 1]$			
Uniform$[0, 1]$	0.503	0.436	0.092
Beta$(2, 1)$	0.730	0.671	0.211
Beta$(0.5, 0.8)$	0.365	0.267	0.057
truncN$(0.5, 0.25)$ on $[0, 1]$	0.524	0.481	0.062
truncN$(0.7, 0.25)$ on $[0, 1]$	0.748	0.690	0.188
Hierarchical: $k_a \mid k_b$ *on* $[0, k_b]$ *(so* $k_a \leq k_b$ *pointwise)*			
$k_b \sim$ truncN$(0.65, 0.25)$ on $[0, 1]$, k_a trunc. on $[0, k_b]$	0.736	0.292	0.006
$k_b \sim$ truncN$(0.25, 0.15)$ on $[0, 1/2]$, k_a trunc. on $[0, k_b]$	0.020	0.002	0.000
Empirical (Part 2)	**0.725**	**0.583**	**0.300**

C Instructions (English translation)

C.1 Random dictator

General Instructions

Welcome to the experiment.

All participants receive the same instructions. Please read them carefully.

During the experiment, you are not permitted to communicate with any other participants, and you must turn off your cell phone. If you have any questions, please raise your hand, and the experimenter will come to you to answer your question privately.

You will receive 5 Euros for showing up. The amount of money you earn on top of this depends on decisions made by you and other participants. Earnings are expressed in points during the experiment. Points convert to Euros in the following way: 10 points = 1 Euro. Your payment will be made privately at the end of the experiment.

The experiment is anonymous. Your identity will not be disclosed to other participants, and the identity of others will not be disclosed to you.

The experiment consists of three parts. You will receive separate instructions for each part. The instructions for Part 1 are on the next page. Instructions for Parts 2 and 3 will be provided once the previous part has been completed.

Instructions for Part 1

At the beginning of Part 1, all participants will be randomly divided into groups of six members each. This means that you will be assigned to a group composed of you and five other people randomly selected by the computer. You will remain in the same group for the entire duration of Part 1.

You will then participate in five consecutive interactions. In each interaction, you will be paired with a different member of your group.

Rules of the Interactions

The structure of all five interactions will be the same. In each interaction, you must choose between two actions: action Y or action Z. You will make your choice without knowing the action chosen by the person you are paired with, and they will not know your choice either.

The points you earn in each interaction will depend on both your chosen action and the action of the person you are paired with.

The table displayed on the screen shows the points you and the person you are paired with can earn in any of the five interactions, depending on both your choice and their choice.

This earnings table will remain the same throughout all five interactions in Part 1.

At the end of each interaction, you will get to see the choice of the paired participant and your earnings in points in that interaction.

The points you earn across the five interactions will be summed up and converted to Euros at the end of the experiment.

Summary of Part 1

Participants will be randomly divided into groups of six.

You will participate in five interactions.

In each interaction, you will be paired with a different member of your group.

You will be paid based on the points earned in all five interactions.

Control Questions

Before the decision-making in Part 1 begins, you will be asked to answer a series of control questions on your computer screen. Once everyone has answered all questions correctly, Part 1 will begin.

Instructions for Part 2

At the beginning of Part 2, all participants will once again be randomly divided into new groups of six. This means that you will be assigned to a group composed of you and five other people randomly selected by the computer. You will remain in the same group for the entire duration of Part 2.

You will then participate in another five consecutive interactions. In each interaction, you will be paired with a different member of your group.

Rules of the Interactions

The five interactions in Part 2 will have the same structure as those in Part 1. However, the earnings may differ. Some groups may use the same earnings table as in Part 1 (referred to as the "original earnings table"), while others may use a different table (referred to as the "alternative earnings table").

Your screen displays both the original and alternative earnings tables, showing the points you and the person you are paired with can earn in any of the five interactions, depending on both your choice and their choice.

At the end of each interaction, you will get to see the choice of the paired participant and your earnings in points in that interaction.

The points you earn across the five interactions of Part 2 will also be summed up and converted to Euros at the end of the experiment.

Assignment of the Earnings Table

The assignment of the earnings table to your group occurs as follows.

At the beginning of Part 2, each member of your new group will be asked to choose between the original earnings table and the alternative one. The computer will then randomly select one of the six group members, and the table chosen by that person will be used by the entire group. You will be informed of which earnings table your group will use before the interactions begin.

This earnings table will remain the same throughout all five interactions in Part 2.

Before informing you which earnings table your group will use, we will ask you to share your expectations about the choices of two individuals who participated in a similar experiment a few days ago. After Part 1, one person was in a group that the computer randomly assigned to the original earnings table, while the other was in a different group assigned to the alternative earnings table. You will be rewarded based on accuracy of your expectations.

Summary of Part 2

Participants will again be randomly divided into groups of six.

You will choose between the original earnings table (the same as in Part 1) and the alternative one.

Your group will use the earnings table chosen by one of the six group members, randomly selected by the computer.

Before knowing which earnings table your group will use, you will be asked to share your expectations about how others behaved under both earnings tables.

You will participate in five interactions.

In each interaction, you will be paired with a different member of your group.

You will be paid based on the points earned in all five interactions.

Control Questions

Before the decision-making in Part 2 begins, you will be asked to answer a series of control questions on your computer screen. Once everyone has answered all questions correctly, Part 2 will begin.

Instructions for Part 3

The structure of Part 3 is identical to Part 2, except there will be only one interaction in Part 3:

All participants will once again be randomly divided into new groups of six.

You will be paired with one member of your group to participate in one additional interaction.

The additional interaction will have the same structure as those in the previous parts.

Your group will use either the original or the alternative earnings table, based on the same procedure as in Part 2.

You will be asked to share your expectations about the behavior of other individuals (this time, there will be no reward for accuracy).

The points earned in the interaction in Part 3 will also be converted into Euros at the end of the experiment.

After completing Part 3, you will be asked to fill out a short questionnaire.

C.2 Coin flip

General Instructions

Welcome to the experiment.

All participants receive the same instructions. Please read them carefully.

During the experiment, you are not permitted to communicate with any other participants, and you must turn off your cell phone. If you have any questions, please raise your hand, and the experimenter will come to you to answer your question privately.

You will receive 5 Euros for showing up. The amount of money you earn on top of this depends on decisions made by you and other participants. Earnings are expressed in points during the experiment. Points convert to Euros in the following way: 10 points = 1 Euro. Your payment will be made privately at the end of the experiment.

The experiment is anonymous. Your identity will not be disclosed to other participants, and the identity of others will not be disclosed to you.

The experiment consists of three parts. You will receive separate instructions for each part. The instructions for Part 1 are on the next page. Instructions for Parts 2 and 3 will be provided once the previous part has been completed.

Instructions for Part 1

At the beginning of Part 1, all participants will be randomly divided into groups of six members each. This means that you will be assigned to a group composed of you and five other people randomly selected by the computer. You will remain in the same group for the entire duration of Part 1.

You will then participate in five consecutive interactions. In each interaction, you will be paired with a different member of your group.

Rules of the Interactions

The structure of all five interactions will be the same. In each interaction, you must choose between two actions: action Y or action Z. You will make your choice without knowing the action chosen by the person you are paired with, and they will not know your choice either.

The points you earn in each interaction will depend on both your chosen action and the action of the person you are paired with.

The table displayed on the screen shows the points you and the person you are paired with can earn in any of the five interactions, depending on both your choice and their choice.

This earnings table will remain the same throughout all five interactions in Part 1.

At the end of each interaction, you will get to see the choice of the paired participant and your earnings in points in that interaction.

The points you earn across the five interactions will be summed up and converted to Euros at the end of the experiment.

Summary of Part 1

Participants will be randomly divided into groups of six.

You will participate in five interactions.

In each interaction, you will be paired with a different member of your group.

You will be paid based on the points earned in all five interactions.

Control Questions

Before the decision-making in Part 1 begins, you will be asked to answer a series of control questions on your computer screen. Once everyone has answered all questions correctly, Part 1 will begin.

Instructions for Part 2

At the beginning of Part 2, all participants will once again be randomly divided into new groups of six. This means that you will be assigned to a group composed of you and five other people randomly selected by the computer. You will remain in the same group for the entire duration of Part 2.

You will then participate in another five consecutive interactions. In each interaction, you will be paired with a different member of your group.

Rules of the Interactions

The five interactions in Part 2 will have the same structure as those in Part 1. However, the earnings may differ. Some groups may use the same earnings table as in Part 1 (referred to as the "original earnings table"), while others may use a different table (referred to as the "alternative earnings table").

Your screen displays both the original and alternative earnings tables, showing the points you and the person you are paired with can earn in any of the five interactions, depending on both your choice and their choice.

At the end of each interaction, you will get to see the choice of the paired participant and your earnings in points in that interaction.

The points you earn across the five interactions of Part 2 will be also summed up and converted to Euros at the end of the experiment.

Assignment of the Earnings Table

The assignment of the earnings table to your group occurs as follows.

At the beginning of Part 2, the computer will randomly assign your new group to either the original earnings table or the alternative one. You will be informed of which earnings table your group will use before the interactions begin.

This earnings table will remain the same throughout all five interactions in Part 2.

Before informing you which earnings table your group will use, we will ask you to share your expectations about the choices of two individuals who participated in a similar experiment a few days ago. After Part 1, one person was in a group that the computer randomly assigned to the original earnings table, while the other was in a different group assigned to the alternative earnings table. You will be rewarded based on accuracy of your expectations.

Summary of Part 2

Participants will again be randomly divided into groups of six.

The computer will randomly assign either the original earnings table (the same as in

Part 1) or the alternative earnings table to your group.

Before knowing which earnings table your group will use, you will be asked to share your expectations about how others behaved under both earnings tables.

You will participate in five interactions.

In each interaction, you will be paired with a different member of your group.

You will be paid based on the points earned in all five interactions.

Control Questions

Before the decision-making in Part 2 begins, you will be asked to answer a series of control questions on your computer screen. Once everyone has answered all questions correctly, Part 2 will begin.

Instructions for Part 3

The structure of Part 3 is identical to Part 2, except there will be only one interaction in Part 3:

All participants will once again be randomly divided into new groups of six.

You will be paired with one member of your group to participate in one additional interaction.

The additional interaction will have the same structure as those in the previous parts.

Your group will use either the original or the alternative earnings table, based on the same procedure as in Part 2.

You will be asked to share your expectations about the behavior of other individuals (this time, there will be no reward for accuracy).

The points earned in the interaction in Part 3 will also be converted into Euros at the end of the experiment.

After completing Part 3, you will be asked to fill out a short questionnaire.

D Additional tables

Table D.1: Treatment effects: Cooperation in Part 1

	(1) Rounds 1–5	(2) Round 5
PriorityTax	0.015	0.006
	(0.047)	(0.066)
Priority	-0.078*	-0.093*
	(0.042)	(0.054)
Controls	✓	✓
Mean dep. var. (Tax)	0.308	0.183
Wald p: PriorityTax = Priority	0.082	0.148
Observations	1,800	360
Matching groups	60	60

Notes: Data from Part 1 only. Estimates are from linear probability models of cooperation decisions. Tax treatment is the reference category. Column (1) uses all Part 1 rounds; Column (2) restricts to Round 5. Standard errors (in parentheses) are clustered at the Part-1 matching group. All specifications include the control set indicated in the table. Wald tests of equality of the corresponding treatment coefficients are reported in the lower panel. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table D.2: Treatment effects: Support for bad policy in Part 2 (Additional controls)

	Full sample		Below-median mistakes	
Dep.var.: Vote for bad policy	(1)	(2)	(3)	(4)
PriorityTax	-0.123*	-0.114*	-0.180**	-0.177**
	(0.068)	(0.067)	(0.078)	(0.076)
Priority	-0.440***	-0.442***	-0.513***	-0.504***
	(0.049)	(0.056)	(0.066)	(0.076)
Right leaning	0.020**	0.028***	0.033***	0.038***
	(0.010)	(0.010)	(0.012)	(0.012)
Guessing game choice	-0.001	-0.001	-0.001	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)
Controls		✓		✓
Mean dep. var. (Tax)	0.710	0.710	0.785	0.785
Wald p: PriorityTax = Priority	0.000	0.000	0.000	0.001
Observations	320	320	179	179
Matching groups	60	60	59	59

Notes: Data from Part 2 only. Estimates are from linear probability models. Standard errors (in parentheses) are clustered at the Part-1 matching group. Controls include the number of quiz attempts, age, gender, occupation, education level and field, and high-school grade. Columns (3)–(4) exclude participants whose number of mistakes in answering the quiz on the instructions exceeds the median (computed among non-control sessions); the number of quiz mistakes is absorbed by the sample restriction. The lower number of observations relative to Table 4 is due to missing answers to the political orientation question. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table D.3: Treatment effects: Support for bad policy in Part 2 (Probit)

	Full sample		Below-median mistakes	
Dep.var.: Vote for bad policy	(1)	(2)	(3)	(4)
PriorityTax	-0.142**	-0.129**	-0.179**	-0.171**
	(0.064)	(0.061)	(0.079)	(0.073)
Priority	-0.425***	-0.427***	-0.496***	-0.504***
	(0.046)	(0.050)	(0.061)	(0.065)
Controls		✓		✓
Mean dep. var. (Tax)	0.725	0.725	0.786	0.786
Wald p: PriorityTax = Priority	0.000	0.000	0.000	0.000
Observations	360	360	195	195
Matching groups	60	60	60	60

Notes: Data from Part 2 only. Estimates are average marginal effects from probit models. Standard errors (in parentheses) are clustered at the Part-1 matching group. Controls include the number of quiz attempts, age, gender, occupation, education level and field, and high-school grade. Columns (3)–(4) exclude participants whose number of mistakes in answering the quiz on the instructions exceeds the median (computed among non-control sessions); the number of quiz mistakes is absorbed by the sample restriction. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table D.4: Beliefs and support for the bad policy

Dep.var: Vote for bad policy	(1)	(2)
Beliefs: good policy	-0.009***	-0.009***
	(0.001)	(0.001)
Beliefs: bad policy	0.005***	0.005***
	(0.001)	(0.002)
PriorityTax × Beliefs: good policy	0.005***	0.003
	(0.002)	(0.002)
Priority × Beliefs: good policy	0.008***	0.009***
	(0.003)	(0.003)
PriorityTax × Beliefs: bad policy	0.004*	0.002
	(0.002)	(0.002)
Priority × Beliefs: bad policy	0.001	0.001
	(0.003)	(0.003)
Controls		✓
Observations	360	360
Clusters	60	60
Adjusted R^2	0.285	0.295

Notes: Data from Part 2 only. Estimates are from linear probability models. Standard errors (in parentheses) are clustered at the Part-1 matching group. Belief variables are measured on a 0–100 scale, with 0 indicating no expected cooperation and 100 indicating certainty about cooperation. All regressions include treatment dummies. Controls include the number of quiz attempts, age, gender, occupation, education level and field, and high-school grade. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table D.5: Belief differential and support for the bad policy

	Tax	PriorityTax	Priority
	(1)	(2)	(3)
Δ(Belief good - Belief bad)	-0.008***	-0.006***	-0.003**
	(0.001)	(0.001)	(0.001)
Controls	✓	✓	✓
Belief good − bad			
Observations	120	120	120
Clusters	20	20	20
Adjusted R^2	0.239	0.283	-0.011

Notes: Data from Part 2 only. Estimates are from linear probability models. Standard errors (in parentheses) are clustered at the Part-1 matching group. The belief differential is defined as the difference between beliefs about cooperation under the good policy and beliefs about cooperation under the bad policy, both measured on a 0–100 scale. Controls include the number of quiz attempts, age, gender, occupation, education level and field, and high-school grade. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table D.6: Treatment effects: Beliefs in Part 2

	Under bad policy		Under good policy	
Dep.var.: Beliefs about cooperation	(1)	(2)	(3)	(4)
PriorityTax	1.775	1.391	8.200***	9.445***
	(3.480)	(3.313)	(2.782)	(2.779)
Priority	-8.558***	-8.157***	10.433***	10.204***
	(2.986)	(3.024)	(2.505)	(3.084)
Controls		✓		✓
Mean dep. var. (Tax)	40.392	40.392	68.342	68.342
p-value: PriorityTax = Priority	0.002	0.004	0.329	0.737
Observations	360	360	360	360
Matching groups	60	60	60	60

Notes: Data from Part 2 only. Estimates are from OLS regressions. Standard errors (in parentheses) are clustered at the Part-1 matching group. In Columns (1)–(2), the dependent variable is beliefs about cooperation under the bad policy; in Columns (3)–(4), the dependent variable is beliefs about cooperation under the good policy. Controls include the number of quiz attempts, age, gender, occupation, education level and field, and high-school grade. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table D.7: Treatment effects: Support for bad policy in Part 2 (all priority treatments)

	Full sample		Below-median mistakes	
Dep.var.: Vote for bad policy	(1)	(2)	(3)	(4)
PriorityTaxPlus	-0.075	-0.051	-0.177**	-0.184**
	(0.071)	(0.070)	(0.086)	(0.082)
RevPriorityTax	0.025	0.057	-0.050	-0.084
	(0.069)	(0.068)	(0.085)	(0.082)
Priority	-0.283***	-0.291***	-0.317***	-0.382***
	(0.062)	(0.064)	(0.082)	(0.083)
Controls		✓		✓
Observations	480	480	283	283
Matching groups	80	80	80	80

Notes: Data from Part 2 of the treatments including a priority policy. Estimates are from linear probability models. Standard errors (in parentheses) are clustered at the Part-1 matching group. The reference category is the PriorityTax treatment. Controls include the number of quiz attempts, age, gender, occupation, education level and field, and high-school grade. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table D.8: Vote for bad policy in Part 3 among participants who voted for the bad policy in Part 2

Dep.var.: Vote for bad policy	(1) Tax	(2) PriorityTax	(3) Priority
Good policy in Part 2	-0.087	-0.214	-0.508**
	(0.061)	(0.152)	(0.183)
Controls	✓	✓	✓
Observations	87	70	36
Clusters	10	10	10
Adjusted R^2	0.117	0.022	0.253

Notes: Data from Part 3; participants who voted for the bad policy in Part 2 only. Estimates are from linear probability models. Standard errors (in parentheses) are clustered at the matching-group level. Controls include the number of quiz attempts, age, gender, occupation, education level and field, and high-school grade. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table D.9: Belief update

	Voted BP in Part 2				Voted GP in Part 2			
Dep.var.: Beliefs under:	(1) BP	(2) GP	(3) BP	(4) GP	(5) BP	(6) GP	(7) BP	(8) GP
Part 3	1.031	14.561***	-10.547***	-0.053	-1.661	8.295***	-9.582***	1.709
	(2.427)	(2.318)	(1.917)	(2.102)	(1.622)	(1.641)	(2.676)	(1.733)
BP in Part 2	3.568	1.348			0.838	1.779		
	(3.567)	(3.187)			(3.399)	(2.257)		
BP in Part 2 × Part 3	-11.578***	-14.614***			-7.921**	-6.586***		
	(3.116)	(3.164)			(3.413)	(2.328)		
GP in Part 2			-3.568	-1.348			-0.838	-1.779
			(3.567)	(3.187)			(3.399)	(2.257)
GP in Part 2 × Part 3			11.578***	14.614***			7.921**	6.586***
			(3.116)	(3.164)			(3.413)	(2.328)
Part 2-3 update (exp.), p-val.	0.00	0.98	0.67	0.00	0.00	0.33	0.31	0.00
Controls	✓	✓	✓	✓	✓	✓	✓	✓
Observations	386	386	386	386	334	334	334	334
Clusters	30	30	30	30	30	30	30	30
Adjusted R^2	0.100	0.176	0.100	0.176	0.099	0.143	0.099	0.143

Notes: Data are from the Tax, PriorityTax, PriorityTaxPlus, and Priority treatments and are drawn from Parts 2 and 3. The dependent variable is beliefs about cooperation, measured on a 0–100 scale. Estimates come from linear regressions of beliefs on a Part-3 indicator and its interaction with experience in Part 2. Experience in Part 2 is defined as playing the bad or the good policy in Part 2. Columns (1)–(4) restrict the sample to participants who voted for the bad policy in Part 2, while Columns (5)–(8) restrict the sample to participants who voted for the good policy in Part 2. All specifications include treatment fixed effects, as well as the number of quiz attempts, age, gender, occupation, education level and field, and high-school grade. Standard errors (in parentheses) are clustered at the Part-2 matching-group level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

E Additional figures

Figure D.2: Event-study estimates of efficiency

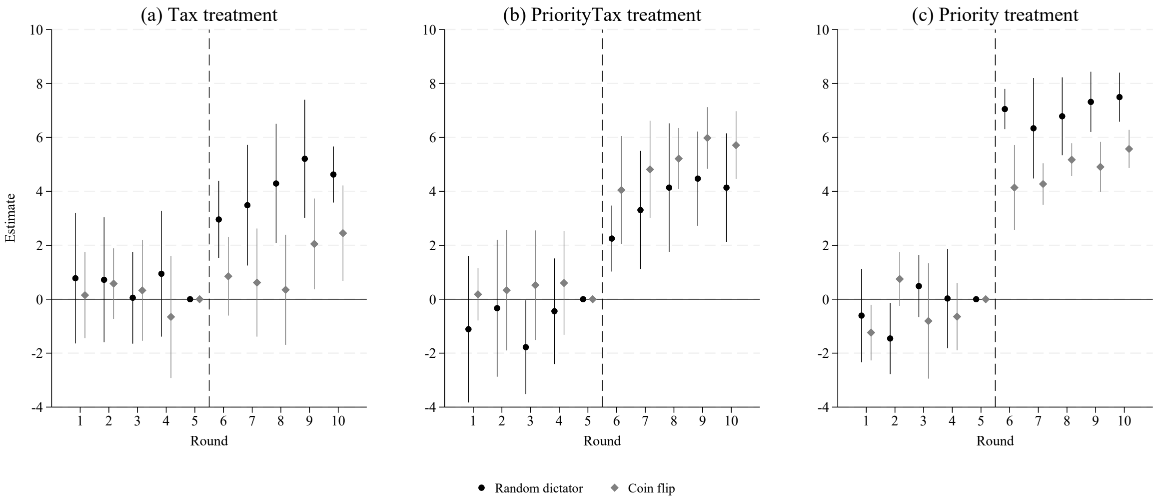

Notes: Each panel reports event-study estimates of efficiency, computed at the pair level as the sum of payoffs. Data from Parts 1 and 2 only. The underlying specification interacts round dummies with an indicator equal to one when the game played in Part 2 is the good policy and zero when participants instead play under the bad policy. Round 5—the final round of Part 1—is the omitted category. Circles denote estimates from the random-dictator conditions; diamonds denote estimates from the corresponding coin-flip control conditions. Standard errors are clustered at the matching-group level.

Figure D.3: Cooperation rates over rounds

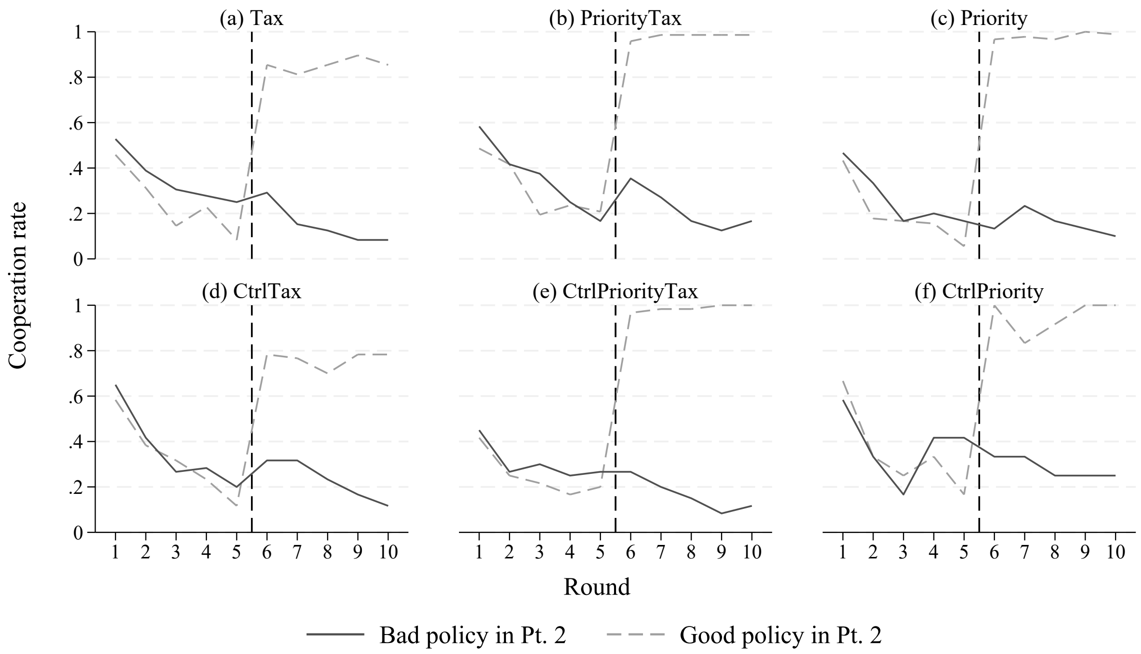

Notes: The figure plots raw cooperation rates in each round of Parts 1 and 2, separately by treatment. Within each treatment, the solid line reports cooperation rates when the game played in Part 2 remains under the bad policy, and the dashed line reports cooperation rates when participants instead play under the good policy. The vertical line at Round 6 marks the start of Part 2.

Figure D.4: Average efficiency over rounds

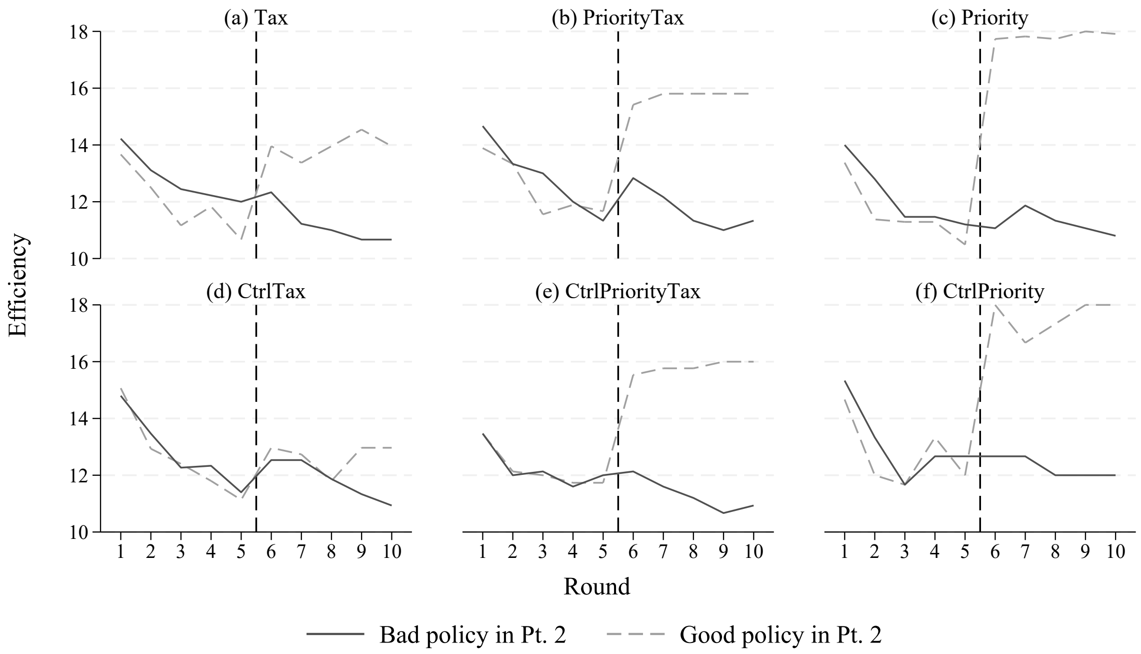

Notes: The figure plots raw mean efficiency in each round of Parts 1 and 2, separately by treatment. Efficiency is computed at the pair level as the sum of payoffs. Within each treatment, the solid line reports efficiency when the game played in Part 2 remains under the bad policy, and the dashed line reports efficiency when participants instead play under the good policy. The vertical line at Round 6 marks the start of Part 2.

Figure D.5: Beliefs under good policy in Part 2

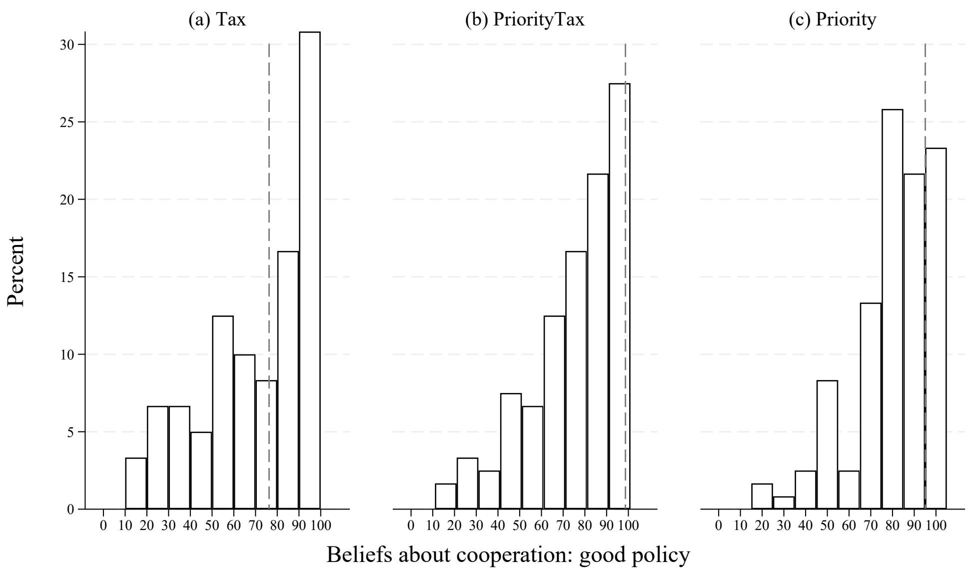

Notes: Each panel shows the distribution of beliefs about cooperation under the good policy in Part 2 by treatment. The histograms report the percentage of participants in 10-point bins of stated beliefs (0–100). The vertical segment marks the actual cooperation rate under the good policy in Part 2 for the corresponding control condition.

Figure D.6: Event-study estimates for PriorityTaxPlus

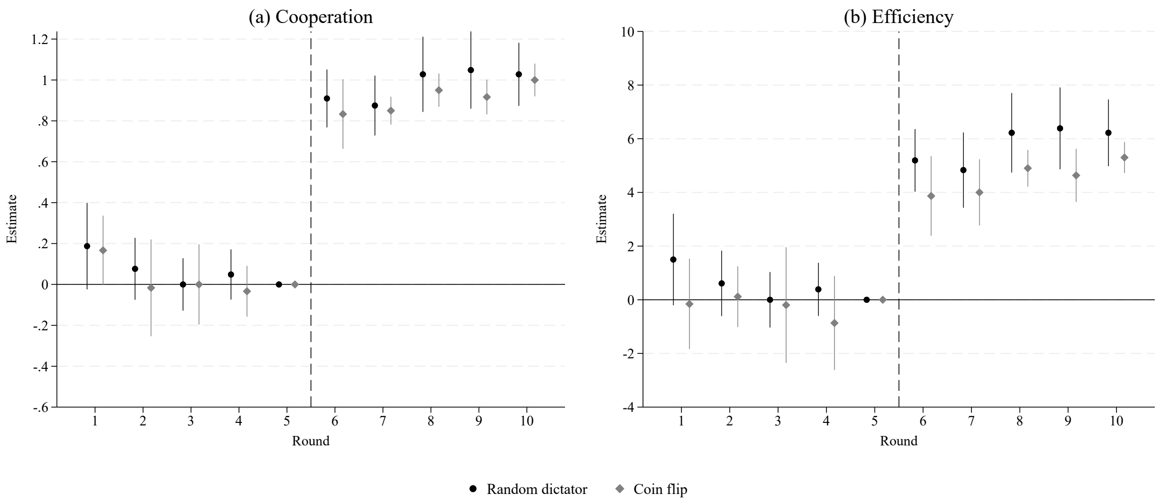

Notes: Each panel reports event-study estimates of cooperation choices (panel a) and efficiency, measured as the sum of pair payoffs (panel b), for the PriorityTaxPlus treatment. Data from Parts 1 and 2 only. The underlying specification interacts round dummies with an indicator equal to one when the game played in Part 2 is the good policy and zero when participants instead play under the bad policy. Round 5—the final round of Part 1—is the omitted category. Circles denote estimates from the random-dictator condition; diamonds denote estimates from the corresponding coin-flip control condition. Estimates come from linear probability models (panel a) and OLS (panel b). Standard errors are clustered at the Part-2 matching-group level.

Figure D.7: Part-1 cooperation, beliefs, and voting

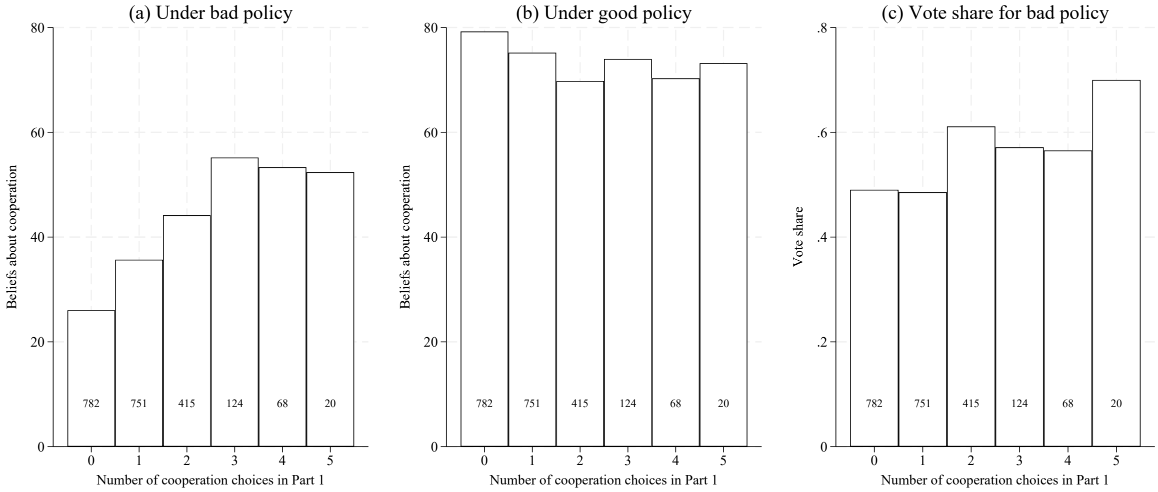

Notes: The figure reports average beliefs about cooperation under bad and good policy (Panels (a) and (b)) and vote shares for the bad policy (Panel (c)) in Part 2, by the number of cooperation choices made in Part 1 (rounds 1–5). Data are pooled across the Tax, PriorityTax, and Priority treatments. Numbers at the base of each bar indicate the number of participants in each bin.

Figure D.8: Event-study estimates for RevPriorityTax

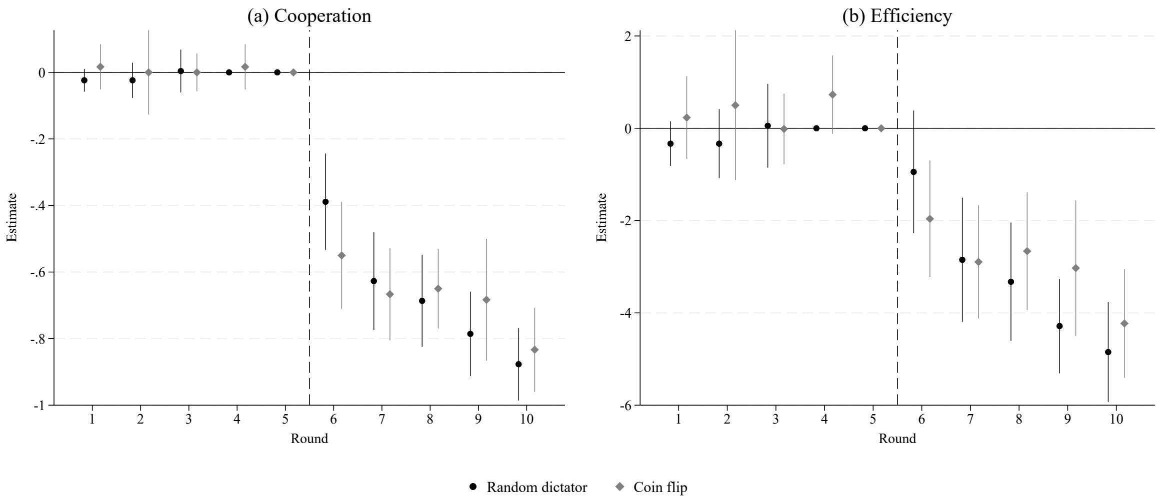

Notes: Each panel reports event-study estimates of cooperation choices (panel a) and efficiency, measured as the sum of pair payoffs (panel b), for the RevPriorityTax treatment. Data from Parts 1 and 2 only. The underlying specification interacts round dummies with an indicator equal to one when the game played in Part 2 is the bad policy and zero when participants instead play under the good policy. Round 5—the final round of Part 1—is the omitted category. Circles denote estimates from the random-dictator condition; diamonds denote estimates from the corresponding coin-flip control condition. Estimates come from linear probability models (panel a) and OLS (panel b). Standard errors are clustered at the Part-2 matching-group level.