Tax, What is it Good for? Perceived Government Efficiency and Redistribution Preferences

Jonathan Kellogg
6,793 words

June 2026

Abstract

A large literature has been devoted to fairness preferences in an effort to explain the gap between predicted and observed levels of redistribution. This paper argues that voters also evaluate the instrument through which redistribution is delivered, and that perceived government efficiency is a determinant of stated redistributive preferences. Building on prospect theory and the "submerged state" by (Mettler, 2011), I develop a framework in which the asymmetric weight voters place on taxation (loss) relative to transfers (gain) is amplified when the state is perceived to convert revenue inefficiently into public goods. I argue that voters use visibility of government projects as a proxy of efficiency. I test this argument using a regression discontinuity design around Britain's highly visible pension fuel payment, exploiting the difference in age of receiving the payment between men and women. I find that receiving a sudden increase of visible government spending increases the willingness to redistribute in Britain, specifically among British women, in line with the predicted gender difference. The findings have two implications. First, conventional estimates of the ideological component of redistributive preferences are biased upward by the omission of perceived efficiency, suggesting that observed cross-national and cross-individual variation in redistributive demand reflects evaluations of state capacity as well as disagreement over inequality itself. Second, the visibility of state action feeds back into the feasible policy set not only through the political mobilization of beneficiaries, as the existing policy-feedback literature emphasizes, but also through its effect on aggregate demand for redistribution. The paper contributes to debates on the microfoundations of redistributive politics and on the downstream political consequences of policy design.

Keywords: redistribution; political economy; policy feedback; government efficiency; Meltzer-Richard

1 Introduction

Critiques and extensions of the canonical model of redistribution by Meltzer and Richard (1981) are nearly as old as the model itself (Alesina and Giuliano, 2011). Rising inequality has not produced the redistributive response that median-voter logic predicts. A substantial literature has sought to resolve this puzzle by enriching the preference structure of the median voter, focusing on social trust, beliefs about fairness and mobility, ethnic heterogeneity, and other-regarding preferences (Rueda and Stegmueller, 2016).

Comparatively little attention has been paid to what voters think about the instrument through which redistribution is delivered. This is a striking omission given that Meltzer and Richard's original formulation explicitly uses the deadweight loss of taxation as part of the voter's calculus. If voters perceive that a pound taxed yields substantially less than a pound transferred, their demand for redistribution should fall even when their distaste for inequality does not. This paper develops and tests the argument that perceived government efficiency is a first-order determinant of redistributive preferences, operating independently of, and potentially confounding, the ideological and identity-based mechanisms that dominate the existing literature. The argument rests on two well-established foundations. First, the demonstration by Kahneman and Tversky (1979) that losses loom larger than equivalent gains. This implies that the perceived cost of taxation will weigh more heavily in voters' calculations than the perceived benefit of transfers, with the asymmetry amplified when voters doubt that taxes are efficiently converted into public goods. This application of prospect theory to political economy follows the example of Passarelli and Del Ponte (2020). Second, Mettler (2011)'s account of the "submerged state" establishes that voters systematically fail to credit government for its outputs. These two insights suggest that the redistributive preferences observed in survey data conflate two distinct judgments: how much inequality voters wish to correct, and how much they trust the state to correct it without substantial waste. To identify the causal effect of perceived efficiency on redistributive preferences, I exploit variation in policy visibility. A line of work running from Pierson (1993) through Soss (1999) to Finkelstein (2009) demonstrates that government interventions are not equally legible to citizens, and that visibility is shaped by administrative architecture in ways that are plausibly exogenous to underlying preferences over redistribution. Visibility thus serves as a tractable proxy for the efficiency signal voters receive, permitting a regression discontinuity design that I describe in Section 4.

The paper makes two contributions. First, it advances the long-running debate over redistributive preferences (Alesina and Giuliano, 2011). This does not falsify existing accounts, but it does suggest that the ideological component of redistributive preferences has been overstated, and that estimates which omit perceived efficiency are likely biased. Second, the paper contributes to the literature on the downstream political effects of policy design (Pierson, 1993; Soss, 1999; Mettler, 2011) by demonstrating that the visibility of state action feeds back into the feasible policy set through its effect on demand for redistribution

itself.

The remainder of the paper proceeds as follows. Section 2 situates the argument in the literature on redistributive preferences and policy feedback. Section 3 presents the theoretical framework. Section 4 describes the data and identification strategy. Section 5 reports results. Section 6 concludes.

2 Literature Review

The argument speaks to three broad literatures: The inequality literature, the literature on downstream political effects of policy choices, and the papers on political behavior conducted by behavioral economists. Below I discuss my intervention to these literature in more depth below.

2.1 Preference-based responses to the redistribution puzzle

As discussed above, much work has been done already to explain why voters diverge in their behavior from the predictions of Meltzer and Richard (1981). Alesina and Giuliano (2011) and more recently, the literature chapter of Cavaillé (2023), offer a good survey of the field: beliefs about social mobility and fairness, identity-based discounting of out-groups, ethnic heterogeneity, and other-regarding preferences each enter as modifications to the voter's objective function. Rueda and Stegmueller (2016) show that altruism toward the poor is a quantitatively significant determinant of redistributive support, conditional on beliefs about whether the poor are deserving. They discuss the condition under which redistribution preferences translate to voting behavior. Dimick et al. (2017) extend this to the upper end of the income distribution specifically. Habyarimana et al. (2007) introduce ethnic heterogeneity as a constraint on cooperative provision, and Starmans et al. (2017) show that opposition to redistribution often reflects preferences for procedural fairness rather than acceptance of inequality.

Yet, this literature is largely mechanism blind. That is to say, they assume at the very least that the divergence between a pound taxed and a pound transferred are not decisive on the preferences of voters. While of course there is only so much papers can focus on, the assumption that stated redistribution preferences purely reflect beliefs around fairness offers a clear opportunity for extension. Especially since there is experimental evidence that voters are efficiency seeking in their redistribution preferences (Almås et al., 2020).

2.2 Policy feedback

A separate literature focuses on the consequences of policy choices for political mobilisation rather than for preferences. According to this logic, redistribution outcomes do not only reflect preferences but are also path dependent on previous policy choices. While it does not directly contradict the research on voter preferences it leaves them aside to focus on the role of policy instruments themselves.

Pierson (1993) establishes that policies create constituencies that defend them, and that the visibility of those policies determines the strength of the resulting feedback loop. Campbell (2003) provides the empirical workhorse for this claim, showing that Social Security recipients are systematically more politically active than comparable cohorts who came of age before the programme, and that the activation is mediated by the programme's legibility. Soss (1999) extends the argument to programme design, showing that administrative legibility teaches citizens different lessons about their own political efficacy. Mettler (2011) catalogues this "submerged" portion of American social policy, the tax expenditures, loan guarantees, and private intermediaries that deliver benefits without registering as government action. Hacker (2002) and Morgan and Campbell (2011) extend this catalogue to the politics of the divided American welfare state and the delegation of social provision to private intermediaries.

The smaller body of work that explicitly focuses on preferences rather than on participation finds the link to be conditional and often muted. Soss and Schram (2007) report that the 1996 welfare reform in the USA produced surprisingly little aggregate movement in opinion despite the magnitude of the change, and Jacobs and Mettler (2018) show that the Affordable Care Act's effect on policy support has been strongly mediated by partisanship and by which features of the law respondents actually notice. The comparative literature is more positive: Larsen (2008) and Kumlin (2004) both argue that the institutional design of welfare provision conditions deservingness perceptions and, through them, attitudes toward redistribution. Michener (2018) adds a cautionary note, showing that the fragmented administration of Medicaid produces uneven and sometimes demobilising effects on its beneficiaries.

Yet, these more positive accounts lack causal rigor and a micro foundation of individual preferences, favoring a description of the impact of a whole scale policy approach. This paper seeks to understand the effect of individual level exposure.

2.3 Loss aversion in fiscal behaviour

A third strand applies prospect theory to fiscal behaviour. The underlying claim that losses loom larger than equivalent gains is canonical (Kahneman and Tversky, 1979), and its application to political choice was established by Quattrone and Tversky (1988), who showed that voters' reference points and the framing of policy alternatives generate systematic departures from the expected-utility benchmark. Individual variation in loss aversion exists, including along demographic lines (Borghans et al., 2009), but is too stable across the population to account for the cross-sectional heterogeneity that motivates this paper.

The fiscal application has developed along two lines. The first treats the tax payment itself as the loss. Engström et al. (2015) show, in the Swedish context, that framing tax obligations as losses adjusts reporting behavior. A parallel literature on tax salience demonstrates that the perceived weight of a tax depends sharply on the channel through which it is collected. Finkelstein (2009) exploits the introduction of electronic toll collection to show that less

salient taxes elicit substantially weaker behavioural responses, and Chetty et al. (2009) establish the same asymmetry for consumer goods when sales taxes are posted inclusive of price. Sausgruber and Tyran (2011) corroborate the pattern experimentally, showing that voters under-perceive indirect taxation and over-tolerate it as a result. The implication is that the loss side of the fiscal ledger is not a fixed psychological quantity; it is itself shaped by institutional design.

Recent research has extended this literature to attitudes toward government benefits and redistribution, including studies by Kishishita and Matsumoto (2024) and Yamamura and Ohtake (2026). Both find that individuals are more supportive of redistribution when they trust the government to spend public funds efficiently and believe that redistribution will benefit ordinary citizens rather than political or economic elites. However, these studies rely on survey experiments that manipulate perceptions of trust rather than examining responses to actual government spending. Moreover, both analyses focus exclusively on the Japanese context. This paper extends their argument to the United Kingdom by leveraging real government spending interventions, thereby addressing concerns about external validity that frequently arise in survey-experimental research (Barabas and Jerit, 2010).

3 The Argument

This paper argues that perceived government efficiency is a key driver of stated redistribution preferences. Because efficiency is hard to observe directly, voters rely on the visibility of government projects as a proxy. Since Meltzer and Richard (1981) predict that richer voters are always less supportive of redistribution regardless of efficiency, the mechanism developed here applies primarily to voters lower down the income scale.

3.1 A voter's utility from redistribution

I take the canonical Meltzer–Richard setting as the starting point and modify it in two ways. Consider a voter i with pre-tax income y_i, facing a flat tax rate τ and receiving a lump-sum transfer financed by mean tax revenue $\tau\bar{y}$. The voter's utility from the tax–transfer bundle is

$$U_i(\tau, \theta, \lambda) = \theta\,\tau\bar{y} \;-\; \lambda\,\tau y_i, \tag{1}$$

where $\theta \in [0, 1]$ is the voter's belief about the fraction of each pound collected that reaches citizens as benefits, and $\lambda \geq 1$ is the loss-aversion coefficient on the tax payment. The condition for a below-mean voter to support redistribution is

$$\theta \;>\; \lambda\,\frac{y_i}{\bar{y}}. \tag{2}$$

Equation 1 nests the standard model as the special case $\theta = 1$, $\lambda = 1$. Two features should be noted up front. First, the utility specification is linear in τ and

therefore characterises the voter's binary support–oppose decision through the threshold condition in equation 2 rather than delivering a continuous Meltzer–Richard tax-rate optimum; the deadweight-loss term that pins down an interior τ^* in the canonical model is suppressed here because the empirical strategy below identifies movement in stated support for redistribution, not in a chosen tax rate. Second, the remainder of this section unpacks the two parameters that depart from the benchmark, explains why λ is taken as given while θ is the object of interest, and characterises the visibility channel through which voters infer θ.

3.2 The three costs of taxation

The parameters in equation 1 are designed to separate three analytically distinct costs that voters attach to taxation:

1. the inherent economic inefficiency created by taxation;
2. the intrinsic disutility of paying taxes, independent of how the proceeds are used;
3. the belief that tax revenues will not be spent effectively.

The first is the deadweight loss. Taxation distorts labour supply, savings, and consumption decisions, so the social cost of raising a pound of revenue exceeds the pound itself. I leave this component to the macroeconomic literature that has measured and modelled it at length (Hauner and Kyobe, 2010; Feldstein, 1999; Hausman, 1981). Since voters within a country face the same fiscal architecture, variation in this component cannot account for the cross-individual heterogeneity in redistributive preferences that motivates the paper.

The second component is captured by λ. Taxpayers code their obligations as losses rather than as forgone gains (Engström et al., 2015; Kahneman and Tversky, 1979), and Passarelli and Del Ponte (2020) formalise the political-economy implication: the asymmetric weight on the tax side of the ledger discourages redistribution even among net beneficiaries. Individual variation in loss aversion is documented (Borghans et al., 2009), but it is sufficiently stable across the population that it cannot, on its own, generate the cross-sectional dispersion in redistributive support observed in survey data. The paper therefore takes λ as given rather than as something to be explained: it sits in equation 1 to ensure that the tax payment and the transfer are evaluated on the correct psychological scales, but it is not the source of the variation the empirical strategy exploits.

The third component is captured by θ, and is the object of interest. Even a non-corrupt state may convert revenue into public goods at a poor rate. An important contribution to this literature is provided by Yamamura and Ohtake (2026), who argue that trust in government redistribution rests on two distinct dimensions: perceived competence and perceived benevolence. In other words, citizens are more likely to support redistributive policies when they believe that

the government is capable of spending public funds effectively and that it acts in the interests of ordinary citizens rather than being captured by elites.

From the perspective of voters, however, this distinction may be less meaningful in practice. Citizens often observe the outcomes of government action but may find it difficult to disentangle whether policy failures result from a lack of capacity or a lack of benevolent intent. Since this paper focuses on the visibility of government spending as a signal of government efficiency, I do not attempt to separate these two mechanisms. Highly visible government spending can plausibly signal both competence and benevolence, as it demonstrates the state's capacity to deliver services while also making the benefits of public expenditure more apparent to citizens. Consequently, distinguishing between these dimensions is not central to the argument advanced in this paper.

3.3 The visibility signal

The inference problem voters face resembles a principal–agent problem. The voter is the principal. Since voter is not aware of the exact value of dead-weight loss and thus how much taxes are optimal, they need some information to make their decision. The agent's type, captured by θ, is the rate at which it transforms a pound of revenue into benefits that reach citizens. The principal cares about θ but does not observe it. Public administration is opaque, and individual exposure to the state is limited to specific encounters. Voters therefore must infer θ from a thin set of observable outputs that stand in for the unobservable underlying type.

The visibility literature reviewed in Section 2.2 establishes the empirical pattern of how that inference proceeds. When the outputs of the state are legible, they function as a credible signal of the agent's type, and citizens credit them. When they are submerged by diffuse public services, the signal is degraded, and citizens do not. In the spirit of Holmstrom and Milgrom (1991), voters can only reward or punish what they can observe. Beliefs about government competence, θ, are therefore updated primarily on visible dimensions of state activity. Institutionally legible programmes generate the signals necessary for electoral accountability, while less visible forms of public provision remain largely outside voters' evaluative framework.

Two structural features of the welfare state push the inferred θ downward. First, the costs of taxation are typically delivered directly, through payroll deductions and VAT, while the benefits are diffuse and often institutionally invisible. The cost side of the principal–agent contract is fully observable to the principal; the benefit side is only partially observable. Voters therefore update on costs more readily than on benefits, and the implicit monitoring technology over-weights the inputs the agent extracts relative to the outputs it delivers. Second, a large share of social spending in advanced democracies is delivered indirectly, through tax expenditures, contracted-out providers, and private intermediaries that obscure the principal's view of the agent's activity. The agent is partially shielded from the principal's gaze on the output side.

The implication is that the default direction of bias in θ is downward even

when the agent is, in fact, honest and competent, because the monitoring technology available to the principal is asymmetric rather than because the agent is shirking. Exogenous variation in visibility, which holds the underlying generosity of redistribution fixed but improves the signal the principal receives, therefore generates variation in θ and, through equation 1, in redistributive support. The visibility channel is best understood not as a change in what the state does but as a change in what the principal is able to observe of what the state does.

3.4 Hypotheses

The framework generates two testable predictions.

H1 (visibility). Exogenous increases in the visibility of government outputs raise perceived θ and, through equation 1, stated support for redistribution. This is the prediction that the identification strategy in Section 4 is designed to test.

H2 (interaction). The effect of θ on support for redistribution should be strongest among low-income voters. From equation 1, voters support redistribution when $\theta > \lambda(y_i/\bar{y})$. For high-income voters, where $y_i/\bar{y} > 1$, this threshold exceeds the feasible range of θ. As a result, these voters oppose redistribution regardless of how efficient they perceive the state to be, making changes in θ unlikely to affect their preferences.

For low-income voters, however, $y_i/\bar{y} < 1$, so the threshold falls within the feasible range of θ. In this case, even modest increases in perceived government efficiency can move voters from opposing to supporting redistribution. Visibility-driven updates in θ should therefore have their largest effect among those who stand to benefit most from redistribution.

Intuitively, the rich are already opposed to taxation on income grounds, so efficiency trade-offs are not pivotal for them. The poor, by contrast, stand to gain from redistribution if and only if the state converts revenue efficiently enough to clear the loss-aversion-adjusted threshold on their own income.

The empirical implication for Section 5 is that the estimated jump at the visibility cutoff should be larger for respondents with lower lifetime income or lower occupational class.

4 Methodology

4.1 Identification: a visibility cutoff as a regression discontinuity

The argument in Section 3 implies that stated redistributive preferences respond to the visibility of government outputs. Voters do not observe the efficiency parameter θ directly; they infer it from concrete encounters with the state, and those encounters are not equally legible across individuals (Pierson, 1993; Soss, 1999; Mettler, 2011; Finkelstein, 2009).

I exploit an administrative cutoff at which a labelled, separately identifiable government transfer turns on. Around the cutoff, a small change in the assignment variable produces a discrete change in whether the individual receives an institutionally legible signal of state activity, while underlying preferences over inequality, identity, and ideology are smooth. Under standard regularity conditions, the local discontinuity in stated support identifies the causal effect of the visibility cue on θ. The theory predicts that this local effect is positive.

The implementation is the United Kingdom State Pension Age (SPA) cutoff, observed in the 2005 British Social Attitudes Survey (BSA). At an individual's SPA, the Department for Work and Pensions begins a recurring, separately labelled four-weekly transfer, accompanied by a formal award letter and, the following November, the Winter Fuel Payment as a labelled lump sum. The cutoff is administratively automatic and the transfer is the paradigmatic "legible" benefit in the sense of Mettler (2011). In 2005, the State Pension Age (SPA) was 60 for women and 65 for men. Pension eligibility thresholds are often viewed as problematic regression discontinuity cutoffs because they coincide with major life-course transitions, particularly retirement, that may independently affect political attitudes and behaviour.

In the British case, however, this concern is less severe for women. While men tend to retire around the age of 65, women in the UK often remain in the labour force for longer after reaching the SPA (Loretto and Vickerstaff, 2013). As a result, crossing the eligibility threshold at age 60 is less tightly linked to changes in employment status and lifestyle for women than for men. This makes the regression discontinuity design based on women's SPA comparatively cleaner, as treatment assignment is less likely to be confounded by contemporaneous retirement decisions.

The outcome is the canonical BSA item `REDISTRB`, asking whether the government should redistribute income from the better-off to those less well-off, on a five-point Likert scale, reverse-coded so higher values denote stronger support.

A key objection to this identification strategy is that receiving the Winter Fuel Payment changes an individual's income and may therefore affect redistributive preferences through the standard Meltzer–Richard mechanism rather than through visibility. While receipt of the payment does increase income, its use as a treatment remains justified for three reasons.

First, the Winter Fuel Payment is a lump-sum transfer. As such, it closely resembles the lump-sum redistribution assumed in the Meltzer–Richard framework. Because eligible recipients receive the same payment regardless of income, the transfer has only a limited effect on individuals' relative positions within the income distribution. Since the Meltzer-Richard model is focused on relative wealth, it is unlikely to substantially alter preferences through changes in relative economic standing alone.

Second, the payment is both predictable and guaranteed. Individuals approaching eligibility can reasonably anticipate receiving it, meaning that the difference in expected lifetime resources between someone just below and just above the eligibility threshold is small. If individuals incorporate expected future income into their economic decisions, then much of the payment's value

should already be reflected in their preferences before receipt. While individuals may not behave exactly according to a permanent-income model, there is evidence that people adjust behaviour in anticipation of future income streams (Hall, 1978). The primary change induced by crossing the eligibility threshold is therefore not a substantial increase in expected wealth, but rather the salience and visibility of a government transfer. Accordingly, the treatment is best understood as making government redistribution more visible rather than materially altering recipients' economic position.

Finally, it is important to note that this potential confounder would bias the estimates against the proposed theory rather than in its favour. Under standard models of redistribution, an increase in income should make individuals less supportive of redistributive policies, as they stand to contribute more and benefit less from redistribution. By contrast, the theory advanced here predicts that exposure to a visible government transfer increases support for redistribution.

Consequently, if receipt of the payment raises both income and support for redistribution, the income effect works in the opposite direction of the hypothesised visibility effect. Any positive relationship between treatment and support for redistribution therefore cannot easily be explained by the income increase alone. If anything, the estimated effect is likely to understate the true impact of visibility, since the income channel should attenuate rather than amplify the observed relationship.

4.2 Estimator

The quantity of interest is how much stated support for redistribution changes when an individual crosses the SPA cutoff at which the visibility cue switches on. I estimate this with a standard local-linear regression (Calonico et al., 2014),

$$Y_i = \alpha + \tau \cdot \mathbb{1}[X_i \geq c] + \beta_1(X_i - c) + \beta_2(X_i - c) \cdot \mathbb{1}[X_i \geq c] + \varepsilon_i, \qquad (3)$$

fitted by weighted least squares on observations within the MSE-optimal bandwidth of Calonico et al. (2014), using a triangular kernel that places more weight on observations nearer the cutoff. The coefficient τ is the jump in Y at $X = c$, net of the local slopes on each side. I report bias-corrected, robust standard errors following the same authors.

Age in BSA is recorded only in integer years, so the running variable takes a small number of repeated values. This raises two issues. The first is mechanical: `rdrobust` prints a discrete-running-variable warning, which I acknowledge but cannot resolve in the available data. The second is substantive: at exactly $X_i = c$ the cutoff may coincide with other developments that are not part of the visibility effect. In this case, retirement, benefit filing, attitude updating. I therefore report a donut specification that drops the observations at $X_i = c$ and treat it as the preferred estimate, following Lee and Card (2008) and Kolesár and Rothe (2018).

The loss-aversion parameter λ of equation (2) is, by construction, held fixed across the cutoff: nothing about the individual's psychology changes when she

crosses an administrative threshold. Underlying preferences over inequality are assumed smooth, so any local jump in stated support is attributable to the visibility cue rather than to a discrete change in taste for redistribution.

4.3 Identifying assumptions and tests

No manipulation. The running variable is administrative and exogenous: an individual cannot move her date of birth. The Cattaneo et al. (2020) robust manipulation test returns $T = -0.94$ ($p = 0.345$) and therefore does not reject the no-manipulation null. Any imbalance in the density across the cutoff is consistent with the expected demographic pattern of cohort-size variation and differential mortality rather than with strategic behaviour.

Smoothness of potential outcomes. Identification requires that, conditional on observables, the potential distribution of redistributive preferences be smooth across the cutoff. The classical critique of age-based RDs is that retirement, health, and peer-network composition also change at the SPA (Card et al., 2008). I address this in two ways. First, the donut specification dropping the unit mass at the cutoff mitigates discrete retirement-timing effects. Second, the heterogeneity in the results, a significant effect for women at 60, where labour-force participation remains substantial, and a null for men at 65, where retirement is near-universal, is itself a partial test: the cash-and-leisure confounds operate most strongly in the very setting where the effect does not appear.

Placebo cutoffs. I re-estimate equation (3) at shifted cutoffs (± 5 years) at which no policy change occurred. A null at shifted cutoffs is necessary, though not sufficient, for the headline estimate to be read as a policy effect rather than a smooth age-gradient.

Covariate balance. I report local-linear RDs with each baseline covariate as the dependent variable. Significant discontinuities in covariates that should be smooth would imply selection or measurement-error problems.

4.4 Heterogeneity proxy

H_2 predicts that the visibility effect concentrates among net beneficiaries of redistribution. I use NS-SEC occupational class (`RCLASSGP`) as the pre-specified primary proxy for net-beneficiary status, splitting at managerial/professional (1–2) versus intermediate/routine/manual and never-worked (3–5 and 8). Total household income (`HHINCOME`) is reported as a secondary check. NS-SEC is preferred because, in the 55–65 age window relevant to the SPA cutoff, `HHINCOME` bundles occupational and private pension receipts that themselves vary in visibility, blurring the structural concept of a net beneficiary in Meltzer and Richard (1981). NS-SEC, by contrast, captures lifetime occupational position and is uncontaminated by this channel.

Table 1: Main local-linear RD estimates: UK State Pension Age (BSA 2005).

	$\hat{\tau}$	95% CI	p	N
Pooled	+0.342	[−0.245, 1.096]	0.213	1,151
Women (cutoff 60)	+0.704	[0.073, 1.812]	**0.034**	638
Men (cutoff 65)	−0.221	[−1.516, 0.609]	0.403	513

Note. Local-linear RD with triangular kernel and MSE-optimal bandwidth (Calonico et al., 2014). Bias-corrected, robust CIs. Outcome scaled so higher values denote stronger support for redistribution. Bold p-values significant at 0.05.

4.5 Data

The British Social Attitudes data are the 2005 cross-section (UK Data Service SN 5618; National Centre for Social Research, 2007), publicly available under standard data-use agreements. Replication code is available here. The analyses use `R` 4.5.2 with the `rdrobust` and `rddensity` packages of Calonico et al. (2017) and Cattaneo et al. (2020).

5 Results

Table 1 collects the main estimates; robustness, secondary outcomes, and heterogeneity tests follow in Tables 2–3.

5.1 Main estimates

Figure 1 plots the binned local averages of `REDISTRB` against age centred at the respondent's State Pension Age, together with local-linear fits on either side. The visual discontinuity at the cutoff is consistent with the headline estimates reported below.

Pooling men and women, the local-linear RD on `REDISTRB` recovers a positive but imprecise estimate of +0.34 on the five-point scale (95% CI [−0.25, 1.10]). The pooled estimate masks substantial heterogeneity. For women, whose SPA in 2005 was sixty, the estimate is $\hat{\tau} = +0.70$ (95% CI [0.07, 1.81], $p = 0.034$). This is roughly half a within-sample standard deviation, in the direction predicted by H_1. For men, whose SPA was sixty-five, the estimate is small, negative, and far from significant ($\hat{\tau} = -0.22$, $p = 0.40$).

The asymmetry is consistent with the mechanism of Section 3. British labour-force participation among women at sixty in 2005 remained substantial: many respondents had not yet retired and were not yet drawing occupational pensions. The visibility shock at the cutoff landed in a relatively "blank" transfer environment, where the new institutionally legible signal of state activity is least confounded by simultaneous retirement, by the substitution of one pension type for another, or by the change in peer-network composition that

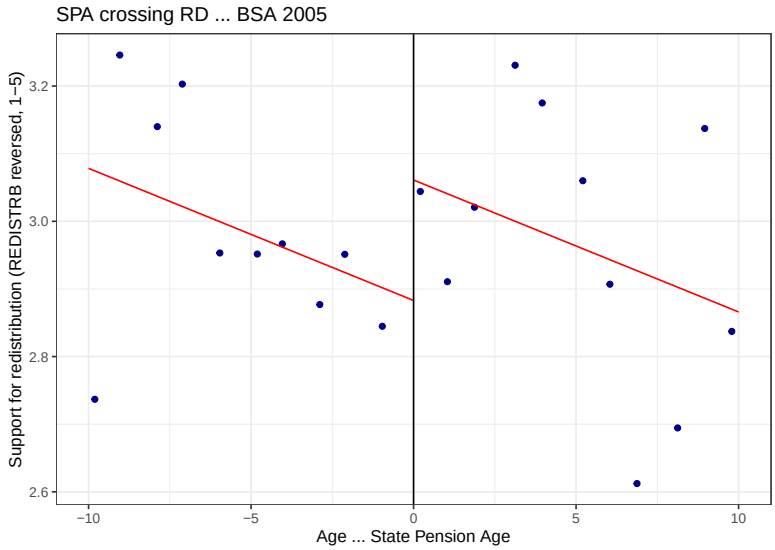

Figure 1: Regression discontinuity at the UK State Pension Age (BSA 2005). Bins are evidence-spaced; lines are local-linear fits with triangular kernel weights.

accompanies full retirement. Among men at sixty-five, by contrast, retirement and occupational-pension receipt are near-universal at the cutoff, and the state pension arrives into an established pattern of pension income. The marginal visibility cue is correspondingly attenuated.

5.2 Robustness

The donut specification, which drops the unit mass at age = SPA to absorb the discrete retirement-timing effect that conventional age-RDs are most criticised for (Card et al., 2008; Kolesár and Rothe, 2018), returns $\hat{\tau} = +0.34$ ($p = 0.44$): smaller and noisier than the women-at-60 estimate, but with the same sign, and consistent with a real but discretely-supported effect in the pooled sample. A global quadratic returns a similar magnitude (+0.45, $p = 0.40$). The placebo cutoffs at SPA ± 5 return +0.15 and −0.28, neither close to significance. The discontinuity in stated redistributive support thus appears only at the policy cutoff.

5.3 Heterogeneity

H_2 predicts that the visibility effect concentrates among net beneficiaries of redistribution. On the pre-specified primary proxy (NS-SEC class, Section 4.4), the result is consistent with H_2 in direction but not in precision: the estimated jump is more than four times larger for lower NS-SEC classes (+0.53)

Table 2: Robustness across specifications (BSA 2005, pooled).

	$\hat{\tau}$	95% CI	N
Main (MSE-optimal h)	+0.342	[−0.245, 1.096]	1,151
Donut RD	+0.337	[−0.633, 1.470]	1,083
Quadratic polynomial	+0.447	[−0.648, 1.621]	1,151
Placebo cutoff, SPA − 5	+0.148	[−0.574, 1.203]	1,151
Placebo cutoff, SPA + 5	−0.283	[−1.233, 0.495]	1,151

Note. Local-linear RD with triangular kernel; MSE-optimal bandwidth except where reported. Bias-corrected, robust CIs.

Table 3: Heterogeneity by income and class (test of H_2, BSA 2005).

Sub-sample	$\hat{\tau}$	95% CI	p	N
Primary split (NS-SEC class):				
Lower NS-SEC class, 3–5 & 8 (pooled)	+0.529	[−0.299, 1.491]	0.192	626
Higher NS-SEC class, 1–2 (pooled)	+0.122	[−0.847, 1.199]	0.736	518
Secondary split (household income):				
Below-median household income (pooled)	+0.038	[−1.316, 0.910]	0.721	520
Above-median household income (pooled)	+0.423	[−0.521, 1.839]	0.273	351
Below-median household income (women)	+0.278	[−1.179, 1.550]	0.790	285
Above-median household income (women)	+0.608	[−0.270, 2.248]	0.124	186

Note. NS-SEC occupational class is the pre-specified primary proxy for net-beneficiary status because it captures lifetime occupational position; household income is reported as a secondary check but is contaminated by occupational and private pension receipts that themselves vary in visibility (see Section 4.5). Bias-corrected, robust CIs. p-values are not adjusted for multiple comparisons.

than for higher (+0.12), neither significant individually. The secondary split by `HHINCOME` runs the opposite way, with the larger estimate in the upper half (+0.42 vs. +0.04 pooled; +0.61 vs. +0.28 for women). As anticipated in Section 4.4, this contamination is plausible: among respondents aged 55–65, those above the income median are disproportionately recipients of already-visible occupational pensions, so the state-pension cue at SPA arrives into a context where the visibility of state income streams is already high, and the marginal cue may operate with comparable strength on both sides of the income median. The NS-SEC result is therefore the more interpretable one and points in the predicted direction, but the evidence is mixed and, since none of the individual estimates reach conventional thresholds of statistical significance, we fail to reject the null hypothesis for H_2. The estimates should also be read against the background of multiple comparisons: with three sub-sample splits and six sub-cells reported here, the measurements are relatively underpowered.

Table 4: Secondary outcomes (BSA 2005).

Outcome	$\hat{\tau}$	95% CI	p	N
`WEALTH` (workers' share, reversed)	+0.027	[−0.679, 0.544]	0.830	1,151
`WELFHELP` (welfare not too generous, rev.)	−0.126	[−1.016, 0.656]	0.673	865
`LEFTRIGH` (left–right scale)	−0.263	[−0.743, 0.262]	0.348	1,153
`TAXSPEND` (tax–spend tradeoff)	−0.039	[−0.635, 0.486]	0.794	682

5.4 Secondary outcomes

Further, I estimate four secondary outcomes with questions that were differently phrased but deal with similar issues. The estimates are tightly clustered around zero. That the headline movement concentrates in the most direct measure of demand for redistribution — "*the government should redistribute income from the better-off to those less well-off*" — and not in adjacent items such as the tax–spend tradeoff or the left–right derived scale is consistent with the visibility channel operating on demand for government redistribution specifically, rather than producing a broad ideological shift. The interpretation is taken up in Section 6.

5.5 Summary

The British design delivers a positive estimate of roughly half a within-sample standard deviation for women at sixty, where the visibility cue is least confounded, and a precise null for men at sixty-five, where it is bundled with the near-universal retirement transition. The effect is concentrated in the most direct redistribution item, survives the donut and quadratic specifications, and disappears at placebo cutoffs. The pattern is consistent with a reading of H_1 in which the visibility channel operates but is detectable where the visibility content of the transfer is high relative to its cash content and the cutoff is not co-located with another large life-course transition. Conversely, the evidence for H_2 is mixed with household income moving in the opposite direction than expected and class in the right direction. However, none of the estimates reach conventional thresholds of statistical significance. We therefore fail to reject the null hypothesis for H_2.

6 Discussion

Overall, the BSA results provide cautious support for the claim that visibility shapes stated redistribution preferences. Although the effect appears to be limited to women, this subgroup is also the least likely to alter behavior along other dimensions, making it the cleanest discontinuity with respect to visibility. The absence of similar effects in the other samples does not necessarily undermine the theory. Instead, these null findings may reflect the broader limitations of age-based regression discontinuity designs, where reaching pension age coincides

with multiple behavioral and financial changes unrelated to the receipt of highly visible benefits.

Firstly, as highlighted above, the gender difference in the BSA sample is in line with the theory that visibility matters but is not overriding of other behaviour. Loretto and Vickerstaff (2013) find that while gender does matter in retirement decisions, those effects are smaller than often anticipated. It is therefore reasonable to assume that, despite becoming eligible for benefits at a younger age, women change their labour-market behaviour only later on. This makes them especially susceptible to the visibility of government action, which is why the causal identification in this group is particularly encouraging. By contrast, as the UK SPA male sample demonstrates, government-efficiency concerns do not appear to be overriding: when they are paired with simultaneous changes in income and behaviour at the cutoff, they are not determinative.

Secondly, the lack of statistically significant effects on other, superficially similar survey questions is also broadly consistent with the proposed interpretation, since many of these questions ask more directly about perceptions of fairness, for example, whether inequality should be reduced. The fact that the same visibility cue does not produce comparable effects across these items suggests that respondents may interpret the redistribution question specifically through a redistribution lens, rather than as a general statement about fairness or inequality.

At the same time, it is important to note that responsiveness to particular survey wording does not always map neatly onto the precise semantics of the question itself (Sjöström and Holst, 2002). Survey instruments are inherently imperfect and can measure attitudes only bluntly. The absence of significant effects on related questions should be interpreted with caution.

The heterogeneity by NS-SEC class lines up with H_2 in the direction predicted, with an estimated effect roughly four times larger for respondents in lower occupational classes than for those in higher classes. Since the BSA cross-section is split twice (by class and by bandwidth), the noise is expected. Notably, the parallel split by total household income contradicts H_2. The divergence between the effects of household income and subjective social status is noteworthy and warrants further investigation. One possible explanation is that higher-income individuals are more likely to have experienced earlier retirement or prior interactions with government benefit programmes, making them more aware of public spending and its benefits. However, the evidence presented here is insufficient to distinguish between these competing mechanisms. As a result, any attempt to provide a definitive theoretical account would be speculative. The findings should therefore be interpreted as suggestive, pointing to an avenue for future research rather than establishing a clear causal mechanism.

An interesting additional question could be posed around how the interaction between visibility and redistributive support varies with the type of policy. The SPA design used in Section 4 bundles two mechanisms: the respondent crosses the cutoff and is simultaneously made more aware of the state and made materially better off by it.

We can imagine a policy space along two dimensions: the visibility of the

intervention to the individual, and the direct material benefit the individual derives from it. Table 5 sets out the four cells. The use of a 2x2 matrix is for simplicity, both factors are more likely scales than binaries.

The SPA, Medicare, and Winter Fuel cutoffs sit in the high-visibility, high-benefit cell. The "submerged state" of Mettler (2011). Tax expenditures, employer-side healthcare subsidies sits in the high-benefit, low-visibility cell. The interesting cell for future work is the high-visibility, low-benefit one: visible public goods that do not directly enrich the respondent, such as a new school for someone without children, or a refurbished hospital ward used by others. These would isolate the visibility channel from the income channel that the present design cannot disentangle.

Table 5: Policy types by visibility and individual material benefit.

	Low visibility	**High visibility**
High individual benefit	*Submerged transfers*: tax expenditures, mortgage interest relief, employer-side healthcare subsidies.	*Labelled transfers*: state pension, Winter Fuel Payment, Medicare. The SPA design.
Low individual benefit	*General administration*: routine bureaucracy, diffuse regulatory activity.	*Visible public goods*: new schools, refurbished hospitals, legible infrastructure used by others.

An adjacent literature on public-good provision and popular support has largely focused on incumbent accountability rather than redistributive demand. Wack (2024) studies electoral accountability when South African voters are exposed to energy blackouts, and Boudot-Reddy and Butler (2024) finds that improved road connectivity in India generates sustained electoral rewards for the ruling party. Both show that visible state outputs shift political behaviour, but they read the response as support for the government in office. An extension to redistributive support would be a natural next step.

7 Conclusion

The canonical Meltzer–Richard (1981) prediction that inequality drives redistribution has given rise to many further refinements of the model. Two literatures are particularly important. The first enriches the preference structure of the median voter, with beliefs over fairness, mobility, identity, and the deservingness of beneficiaries (Alesina and Giuliano, 2011; Rueda and Stegmueller, 2016; Cavaillé, 2023). The second documents the political consequences of how redistribution is delivered, showing that the visibility of state outputs shapes the

political mobilisation of beneficiaries (Pierson, 1993; Soss, 1999; Mettler, 2011).

Neither addresses the question that the original model places at the centre of the voter's calculus: how much of a pound taxed does the voter believe will reach citizens as benefits, and how does that belief enter stated demand for redistribution?

This paper argues that perceived government efficiency is a first-order determinant of redistributive preferences, operating independently of the ideological and identity-based mechanisms that dominate the existing literature and entering the political process upstream of the mobilisation channel that the policy-feedback literature emphasises. I test the argument using regression discontinuities at administrative cutoffs of the UK State Pension Age. The design returns a positive effect on stated support for redistribution of roughly half a within-sample standard deviation. The causal identification of support for redistribution when voters can directly observe government spending represents an important first step in a broader research agenda that places beliefs about government capacity at the centre of political attitudes. Rather than focusing solely on distributive preferences, this approach highlights how citizens' perceptions of the state's ability to deliver public goods and services shape their willingness to support redistribution.

7.1 Limitations

Three limitations deserve explicit acknowledgement. First, the running variable are recorded only in integer years, which produces the mass-point structure documented in Section 4.2 and which weakens the local-linear extrapolation across the cutoff. The donut specifications are reported as the preferred estimate but the underlying problem is one of measurement rather than of identification. A version of BSA with month-of-birth resolution, available under Special Licence at UKDS, would substantively improve precision.

Second, the BSA result rests on a single survey year. The placebo cutoffs discipline the in-design behaviour, but the design cannot replicate across waves on the current data, and a multi-wave version of the design would have to address the gradually-rising women's SPA from 2010 onward. That extension is a natural follow-up but is outside the scope of this paper.

Third, the current analysis relies exclusively on an age-based eligibility threshold. The German Socio-Economic Panel (SOEP) offers substantially greater scope for examining income-based eligibility cutoffs and could therefore provide a useful extension of the research design. Incorporating such cases would strengthen the external validity of the findings by moving beyond pension-related transfers and would help address concerns that other life-course changes coincide with age-based thresholds.

At the same time, income-based discontinuities introduce additional identification challenges. In particular, individuals may strategically adjust their reported income or labour supply in order to remain eligible for benefits. Such behaviour can generate bunching around the eligibility threshold, violating the as-if random assignment assumption that underpins regression discontinuity de-

signs (Kleven, 2016). Any extension using income-based cutoffs would therefore need to carefully test for and account for potential manipulation of the running variable.

7.2 Further Avenues of Research

There are several promising avenues for future research.

First, the distinction between visible public goods and direct individual transfers deserves further attention. While this paper argues that visibility increases support for redistribution, the mechanism may operate differently depending on whether citizens directly receive benefits themselves or merely observe government provision benefiting others. Future research could examine whether visibility has a stronger effect when government spending is personally experienced rather than indirectly observed.

Second, a key limitation of this paper is the lack of easily accessible granular administrative data. Two alternative data sources could help address this issue. The first is the SOEP-Core dataset, which contains substantially more detailed information on income, transfers, and individual characteristics, but is difficult to access because of privacy restrictions. The second is the Danish Forskningsservice, which provides highly granular data on the implementation of cash compensation schemes accompanying green tax increases, alongside measures of redistributive preferences. Such data would make it possible to examine whether visible compensation payments affect support for redistribution in a policy domain beyond pensions. However, access to this dataset is similarly restricted to researchers affiliated with Danish universities.

Although the pension-age regression discontinuity design is well established (Card et al., 2008), it presents challenges because crossing the retirement threshold is associated with multiple simultaneous changes, including increases in cash transfers. The two alternative designs discussed above do not face this same issue, though they may instead be susceptible to strategic behaviour and resulting bunching around the cutoff. Nevertheless, applying these approaches, especially with a donut RD design, could provide a valuable robustness check and strengthen the validity of the paper's central argument.

Third, the framework invites endogenising visibility itself: if perceived efficiency raises redistributive demand, the design of the welfare state is a strategic equilibrium, and the submerged state of Mettler (2011) ceases to be an accident. Fourth, efficiency perceptions plausibility confound other variables that are often seen as influencing ideology, such as ethnic heterogeneity (Alesina and Glaeser, 2004). A more stratified society could isolate people further from the state and thus lead them to believe the state is less good at redistribution rather than just impact their believes of the deservingness of the welfare recipients. Further research that includes efficiency believes could give us more information about the relative power of different redistribution preference determinants.

References

Alesina, A. and Giuliano, P. (2011). Preferences for redistribution. In Benhabib, J., Bisin, A., and Jackson, M. O., editors, *Handbook of Social Economics*, volume 1, pages 93–131. Elsevier.

Alesina, A. and Glaeser, E. L. (2004). *Fighting Poverty in the US and Europe: A World of Difference*. Oxford University Press.

Almås, I., Cappelen, A. W., and Tungodden, B. (2020). Cutthroat capitalism versus cuddly socialism: Are americans more meritocratic and efficiency-seeking than scandinavians? *Journal of Political Economy*, 128(5):1753–1788.

Barabas, J. and Jerit, J. (2010). Are survey experiments externally valid? *American Political Science Review*, 104(2):226–242.

Borghans, L., Heckman, J. J., Golsteyn, B. H. H., and Meijers, H. (2009). Gender differences in risk aversion and ambiguity aversion. *Journal of the European Economic Association*, 7(2–3):649–658.

Boudot-Reddy, C. and Butler, A. (2024). Paving the road to re-election. *Journal of Public Economics*, 239:105228.

Calonico, S., Cattaneo, M. D., Farrell, M. H., and Titiunik, R. (2017). rdrobust: Software for regression-discontinuity designs. *The Stata Journal*, 17(2):372–404.

Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326.

Campbell, A. L. (2003). *How Policies Make Citizens: Senior Political Activism and the American Welfare State*. Princeton University Press.

Card, D., Dobkin, C., and Maestas, N. (2008). The impact of nearly universal insurance coverage on health care utilization: Evidence from Medicare. *American Economic Review*, 98(5):2242–2258.

Cattaneo, M. D., Jansson, M., and Ma, X. (2020). Simple local polynomial density estimators. *Journal of the American Statistical Association*, 115(531):1449–1455.

Cavaillé, C. (2023). *Fair Enough?: Support for Redistribution in the Age of Inequality*. Cambridge University Press.

Chetty, R., Looney, A., and Kroft, K. (2009). Salience and taxation: Theory and evidence. *American Economic Review*, 99(4):1145–1177.

Dimick, M., Rueda, D., and Stegmueller, D. (2017). The altruistic rich? inequality and other-regarding preferences for redistribution. *Quarterly Journal of Political Science*, 11(4):385–439.

Engström, P., Nordblom, K., Ohlsson, H., and Persson, A. (2015). Tax compliance and loss aversion. *American Economic Journal: Economic Policy*, 7(4):132–164.

Feldstein, M. (1999). Tax avoidance and the deadweight loss of the income tax. *Review of Economics and Statistics*, 81(4):674–680.

Finkelstein, A. (2009). E-ztax: Tax salience and tax rates. *Quarterly Journal of Economics*, 124(3):969–1010.

Habyarimana, J., Humphreys, M., Posner, D. N., and Weinstein, J. M. (2007). Why does ethnic diversity undermine public goods provision? *American Political Science Review*, 101(4):709–725.

Hacker, J. S. (2002). *The Divided Welfare State: The Battle over Public and Private Social Benefits in the United States*. Cambridge University Press.

Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy*, 86(6):971–987.

Hauner, D. and Kyobe, A. (2010). Determinants of government efficiency. *World Development*, 38(11):1527–1542.

Hausman, J. A. (1981). Exact consumer's surplus and deadweight loss. *The American Economic Review*, 71(4):662–676.

Holmstrom, B. and Milgrom, P. (1991). Multitask principal–agent analyses: Incentive contracts, asset ownership, and job design. *The Journal of Law, Economics, and Organization*, 7(Special Issue):24–52.

Jacobs, L. R. and Mettler, S. (2018). Policy feedback and public opinion: The politics of the affordable care act. *Perspectives on Politics*, 16(2):345–363.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291.

Kishishita, D. and Matsumoto, T. (2024). Political trust and preferences for redistribution: Wasteful spending and plutocratic influence. SSRN Working Paper No. 4976959.

Kleven, H. J. (2016). Bunching. *Annual Review of Economics*, 8(1):435–464.

Kolesár, M. and Rothe, C. (2018). Inference in regression discontinuity designs with a discrete running variable. *American Economic Review*, 108(8):2277–2304.

Kumlin, S. (2004). *The Personal and the Political: How Personal Welfare State Experiences Affect Political Trust and Ideology*. Palgrave Macmillan.

Larsen, C. A. (2008). *The Institutional Logic of Welfare Attitudes: How Welfare Regimes Influence Public Support*. Ashgate.

Lee, D. S. and Card, D. (2008). Regression discontinuity inference with specification error. *Journal of Econometrics*, 142(2):655–674.

Loretto, W. and Vickerstaff, S. (2013). The domestic and gendered context for retirement. *Human Relations*, 66(1):65–86.

Meltzer, A. H. and Richard, S. F. (1981). A rational theory of the size of government. *Journal of Political Economy*, 89(5):914–927.

Mettler, S. (2011). *The Submerged State: How Invisible Government Policies Undermine American Democracy*. University of Chicago Press, Chicago.

Michener, J. (2018). *Fragmented Democracy: Medicaid, Federalism, and Unequal Politics*. Cambridge University Press.

Morgan, K. J. and Campbell, A. L. (2011). *The Delegated Welfare State: Medicare, Markets, and the Governance of Social Policy*. Oxford University Press.

National Centre for Social Research (2007). British social attitudes survey, 2005. UK Data Service. SN: 5618.

Passarelli, F. and Del Ponte, A. (2020). Prospect theory, loss aversion, and political behavior. Working paper.

Pierson, P. (1993). When effect becomes cause: Policy feedback and political change. *World Politics*, 45(4):595–628.

Quattrone, G. A. and Tversky, A. (1988). Contrasting rational and psychological analyses of political choice. *American Political Science Review*, 82(3):719–736.

Rueda, D. and Stegmueller, D. (2016). *Who Wants What? Redistribution Preferences in Comparative Perspective*. Cambridge University Press, Cambridge.

Sausgruber, R. and Tyran, J.-R. (2011). Are we taxing ourselves? how deliberation and experience shape voting on taxes. *Journal of Public Economics*, 95(1–2):164–176.

Sjöström, O. and Holst, D. (2002). Validity of a questionnaire survey: response patterns in different subgroups and the effect of social desirability. *Acta Odontologica Scandinavica*, 60(3):136–140.

Soss, J. (1999). Lessons of welfare: Policy design, political learning, and political action. *American Political Science Review*, 93(2):363–380.

Soss, J. and Schram, S. F. (2007). A public transformed? welfare reform as policy feedback. *American Political Science Review*, 101(1):111–127.

Starmans, C., Sheskin, M., and Bloom, P. (2017). Why people prefer unequal societies. *Nature Human Behaviour*, 1:0082.

Wack, M. (2024). Does democracy die in darkness? electricity outages & electoral accountability in south africa. SSRN Working Paper.

Yamamura, E. and Ohtake, F. (2026). The widening gap in tax attitudes: Role of government trust in the post covid-19 period. *arXiv preprint arXiv:2603.06098*.