



# Do Means of Program Delivery and Distributional Consequences Affect Policy Support? Experimental Evidence About the Sources of Citizens' Policy Opinions

Vivekinan L. Ashok<sup>1</sup> · Gregory A. Huber<sup>2,3</sup> 

Published online: 5 March 2019

© Springer Science+Business Media, LLC, part of Springer Nature 2019

## Abstract

Recent scholarship argues that citizens' support for specific government programs in the United States is affected by the means through which benefits are delivered as well as the distributional consequences of these policies. In this paper, we extend this literature in two ways through a series of novel survey experiments, deployed on a nationally representative sample. First, we directly examine differences in public support for prospective government spending when manipulating the mode of delivery. Second, we examine whether information about the distributional consequences of two existing government programs affects their popularity. We find that citizens have a preference for indirect spending that is independent of the distributional consequences of a given policy and identify mechanisms that may explain this view. Furthermore, we find little evidence that highlighting the regressive effects of current government programs significantly reduces the demand for their policy benefits. Our findings have implications for understanding the political calculus of policy design and the potential for public persuasion.

**Keywords** Government spending · Public opinion · Policy design

---

**Electronic supplementary material** The online version of this article (<https://doi.org/10.1007/s11109-019-09534-z>) contains supplementary material, which is available to authorized users.

---

✉ Gregory A. Huber  
gregory.huber@yale.edu

Vivekinan L. Ashok  
v.ashok@cornell.edu

<sup>1</sup> Department of Government, Cornell University, 214 White Hall, Ithaca, NY 14853, USA

<sup>2</sup> Department of Political Science, Yale University, New Haven, USA

<sup>3</sup> Institution for Social and Policy Studies, 77 Prospect St., New Haven, CT 06511, USA

Why are some government benefits transferred directly to citizens—in the form of a check or debit card—while others are passed indirectly through the tax code? Does the method of delivery have implications for a program’s survival? Analysts have noted that many redistributive policies, such as welfare and food stamps, are delivered directly to citizens, while others that tend to benefit higher income groups, such as the home mortgage interest deduction, are delivered indirectly.<sup>1</sup> One important theoretical explanation for this difference is that indirect delivery may allow politicians to obscure or “submerge” regressive government spending that would otherwise not survive public scrutiny, while more popular and progressive programs are delivered directly. How policies are delivered therefore has implications for their visibility and political survival.

While this perspective is theoretically intriguing, there are multiple challenges to this account. For one, many of the mostly overtly redistributive programs that are delivered directly, such as cash welfare and public housing, are deeply contested politically and are, at best, only marginally popular. Those popular programs that are delivered directly, such as social security old age benefits, meanwhile, are generally universalistic social insurance programs, rather than overtly redistributive ones. Indeed, many directly delivered programs are structured that way for the practical reason that program beneficiaries either have no tax liability or insufficient liquidity to deliver the benefits via tax rebates or tax reductions. Among currently submerged programs, by contrast, many serve popular goals, such as promoting home ownership. Citizens might therefore still support these programs if they were delivered directly or if they understood more fully the overall distribution of their costs and benefits.<sup>2</sup>

A potentially powerful alternative explanation that we develop in this paper for the preponderance of programs whose benefits are delivered indirectly is that citizens *prefer* indirect to direct delivery. This preference could be the result of several considerations. Recent research on American politics has documented a decline in public confidence towards government institutions and a corresponding rise in antigovernment sentiments.<sup>3</sup> Thus, citizens may prefer indirect delivery of benefits because it is perceived as requiring less government administration, decreasing overall levels of government spending, and allowing fewer opportunities for fraud or misuse.<sup>4</sup> Scholars have also argued that citizens’ views regarding the deservingness

---

<sup>1</sup> See, Hacker and Pierson (2010), Mettler (2011), Campbell (2003), Faricy and Ellis (2013), Haselswerdt and Bartels (2015), among others.

<sup>2</sup> Still another possibility, not addressed in this paper, is that indirectly delivered (regressive) programs could become more popular among program beneficiaries if they were delivered directly because individuals would better understand that they benefited personally. Any change in popularity due to direct program delivery could indeed be smaller or larger than the potentially countervailing effect of the overall distribution of program costs and benefits becoming more transparent.

<sup>3</sup> See, Hetherington (2005), Hetherington and Rudolph (2008), Hacker and Pierson (2017), McCarty et al. (2016), among others.

<sup>4</sup> It is important to note that indirect delivery of government benefits can also be accomplished by using private sector intermediaries. For example, in the case of Medicare, the national health insurance system for seniors, the federal government contracts services and the monitoring of those services to private health care providers. In an extensive examination of Medicare, Morgan and Campbell (2011) argue that

of government beneficiaries influences public support for social spending.<sup>5</sup> As such, it could be the case that citizens view people who receive benefits through tax reductions as taxpayers that expend greater effort to receive benefits than those who get cash transfers. Correspondingly, citizens might see direct delivery as creating future dependence on government assistance.

While some prior scholarship has sought to test both whether citizens prefer direct over indirect program delivery and how program support is affected by knowledge of the distribution of program costs and benefits, results from these studies are mixed and do not probe underlying beliefs that could explain citizens' preferences. We present results from a new series of survey experiments that allow us to isolate citizen preferences about program delivery and test how providing information about who benefits from programs affects their popularity. Overall, we document the existence of strong *primitive* preferences for indirect government spending and suggest explanations for this view. We also show that support for popular indirect government spending, including the Home Mortgage Interest Deduction, remains high even when framed regressively.

In our first experiment, we measure citizen preferences over two hypothetical new policy proposals—involving childhood nutrition and job training—while randomly varying the way in which the government transfers the benefits to eligible individuals. We find that respondents are substantially more supportive of prospective government spending when transfers are made indirectly, via tax expenditures, than when they are direct cash transfers. This preference may be attributable to the fact that citizens report that indirect delivery would cost the government less than direct delivery. In the case of the job training policy, we further find that respondents associate the direct transfer with an increase in their personal tax liability relative to the indirect transfer. With respect to the childhood nutrition policy, respondents view beneficiaries of the indirect transfer as expending more effort to receive benefits compared to citizens who would receive a direct benefit. As a whole, these results suggest that policymakers may actually reflect constituent preferences when they enact policies that use the tax code or other indirect means to deliver policy benefits.

In our second experiment, we measure citizen support for two current government programs—the Home Mortgage Interest Deduction (HMID) and the Unemployment Insurance (UI) system—after randomly manipulating the way the distributional consequences of each policy are framed. Importantly, compared to prior work, we distinguish the effect of providing additional information about these programs from information about their distributional consequences. We find little evidence that framing these policies as having a regressive distributional effect makes them overall unpopular. However, in the case of UI, respondents are more supportive of the policy under the progressive framing relative to the regressive framing. No such difference in support exists when respondents consider the HMID—though we

---

Footnote 4 (continued)

policymakers “delegate” governance due to citizens' divergent preferences for both small government and greater social provision.

<sup>5</sup> See, Applebaum (2001), Gilens (1999), Henry et al. (2004), among others.

do find some suggestive evidence that the regressive framing of the HMID causes respondents to favor modifying the policy in a progressive direction. Taken together, our results show that regressivity alone does not appear to endanger a program's survival.

## Previous Experimental Approaches

Most scholarship that examines the political logic undergirding government spending design has relied on observational data and qualitative methods.<sup>6</sup> However, in the pursuit of more persuasive causal identification, three recent studies employ experimental approaches to investigate the effect of various policy features on public opinion. We review each of these relevant contributions and highlight important questions that remain unanswered.

In the first study, Mettler and Guardino employ a within-subjects survey experimental design that manipulates the type of information citizens receive about the Home Mortgage Interest Deduction (HMID), the Retirement Savings Contribution Tax Credit (RSCTC), and the Earned Income Tax Credit (EITC) (Mettler 2011, Ch. 3). Their key intervention shows respondents two pieces of information: a brief description of the goals of the program and the distribution of benefits by income group for each policy. The treatments frame the HMID and RSCTC as being biased toward higher income individuals and frame the EITC as being biased toward low income individuals. The authors show that this treatment decreases support for the HMID and RSCTC and increases support for the EITC, relative to respondents' baseline views measured at the beginning of the survey. The authors conclude: “[I]f Americans were more informed of who actually benefits from government policies, that opposition to those favoring the affluent would grow and support for those aiding the less well-off would increase, too” (Mettler 2011, p. 59).

Though Mettler and Guardino take an important first step towards identifying the effect of making particular distributional biases visible to citizens, their within-subjects experimental design could lead to several biases that threaten inference.<sup>7</sup> One concern is that there is no untreated control that is also asked policy evaluations twice. A second concern is that it appears a great deal of the change in support for each program comes from differences in rates of “don't know” responses, raising the possibility that latent support (expressed as “don't know” in the pre-treatment baseline) is unchanged.

This concern is addressed in the second study. Faricy and Ellis (2013) replicate and extend results for the HMID and RSCTC using a between-subjects design while adding a new treatment that manipulates the framing of transfers under each policy. Respondents in this experiment are randomly told that benefits under each policy are passed through the tax code or through a cash transfer. Additionally, half the

<sup>6</sup> See, e.g., Arnold (1992), Gilens (2009), Jacoby (1994), Howard (2007), Morgan and Campbell (2011), and Ellis and Stimson (2012).

<sup>7</sup> For a general discussion of the potential biases in within-subject designs, see, Charness et al. (2012).

respondents are shown information about the distribution of benefits that parallel the first study. The authors find additional evidence that information showing each program is biased toward the wealthy reduces support for the HMID and RSCTC relative to the control condition. And, furthermore, they provide preliminary evidence that citizens prefer that current benefits be transferred to households as a tax expenditure rather than as a cash subsidy, although their design is underpowered to reject the null that the means of benefit delivery has no effect on program support for both programs.

Taken together, these two studies show that providing information that frames a policy as having one distributional consequence, in this case advantaging higher (EITC and RSCTC) or lower (EITC) income individuals, alters support for those policies. Respondents appear to oppose more a policy that is biased toward high income individuals and favor more a policy that benefits low income individuals relative to their support when receiving no distributional information. However, it is important to note that the treatments in both studies manipulate two distinct things. Respondents who receive these treatments in either study are being primed to consider the distributional impacts of policies *and* are informed about a particular bias. Given this “bundled” intervention, we are unable to know the extent to which policy support changed simply due to priming distributional concerns rather than the specific information about which income groups benefit more.<sup>8</sup>

Furthermore, both Mettler (2011, Ch. 3) and Faricy and Ellis (2013) examine support only for existing government policies. Respondents’ support in these surveys could reflect some amount of status quo bias (in which individuals have feelings about existing programs due to their background knowledge and evaluate different delivery mechanisms [the manipulation in the second study] in light of their knowledge of how the existing program works), limiting our ability to predict how citizens would respond to new programs that differed in their distributional consequences and means of benefit delivery.

This issue is partially addressed in a third study, which considers two prospective policies—job training for the unemployed and tax breaks to subsidize parental leave from work. Haselswerdt and Bartels (2015) employ a design similar to that of Faricy and Ellis (2013), manipulating the way in which benefits are transferred to households (This study does not present or manipulate distributional consequences). They find that respondents favor tax expenditures over cash transfers when stating support for these new hypothetical programs. This last study is a useful advancement in understanding mass support for new policies, but the authors do not specify or measure the potential mechanisms that underlie respondents’ stated preferences. We do not know why citizens prefer that these new government programs be delivered via the tax code, nor do we know anything about the perceived consequences of using tax expenditures over cash transfers.

To explain why citizens may prefer indirect spending, this prior work focuses on a general dislike for “government” and traceability—concepts that are related to one

---

<sup>8</sup> For a general discussion of the challenges interpreting the effects of bundled treatments, see, Dunning (2012, pp. 300–302).

another. Faricy and Ellis (2013, pp. 56–57) point to evidence that Americans tend to prefer that government be limited in size, and that a general skepticism of government's role in social policy could lead citizens to view indirect spending differently than direct transfers. Building on this argument, Haselswerdt and Bartels (2015 pp. 608–609) theorize that citizens may prefer indirect spending as its effects on the government's budget are less traceable compared to direct spending. Though neither of these studies explicitly measure mechanisms, the arguments put forth are promising and provide a foundation for our inquiry.

We design and field three survey experiments across three independent samples that examine support for both new and existing government programs while manipulating both the means of program delivery and the framing of the distributional consequences of those programs (i.e., the same program is framed progressively or regressively). Importantly, we measure a variety of underlying mechanisms—with respect to cost and perceptions of program beneficiaries—that could explain any differences in support by the type of delivery.

### **Experiment 1: Evidence of Primitive Preference for Indirect Spending**

When politicians decide to spend revenues to subsidize citizen needs, they can often choose between transferring the benefit directly to individuals or indirectly through a tax expenditure. Though both forms of transfers could be viewed as equivalent, researchers have noted that policymakers seem to disproportionately favor tax expenditures over direct transfers when the policy benefits high-income groups more than low-income groups (Mettler 2011; Hacker and Pierson 2010). One explanation for this choice of policy design is that it allows policymakers to obscure regressive policies. However, it could also be the case that tax incentives are not feasible when delivering benefits to individuals who do not pay federal income taxes. Separate from the political logic of program design, it is also possible that citizens themselves prefer indirect government spending, independent of its distributional consequence.

To measure citizen preferences for direct versus indirect government spending, we employ a survey experiment where we randomly manipulate the way in which policy benefits are transferred to individuals. We recruit respondents by contracting the firm Lucid to direct a nationally representative online panel to a survey hosted on the Qualtrics survey platform.<sup>9</sup> Our sample is constructed by combining a general population subsample ( $N = 1550$ , nationally representative with respect to age, gender, race/ethnicity, education-level, and household income) and an income-targeted subsample ( $N = 1797$ , nationally representative with household incomes between \$55,000 and \$100,000).<sup>10</sup> We create post-stratification weights per the 2016 American Community Survey (ACS) to correct for oversampling households with incomes between \$55,000 and \$100,000 and to better estimate the magnitudes of

<sup>9</sup> For information regarding Lucid's specific recruitment method, see: [https://luc.id/wp-content/uploads/2017/07/IRB-Methodology\\_Lucid\\_.pdf](https://luc.id/wp-content/uploads/2017/07/IRB-Methodology_Lucid_.pdf).

<sup>10</sup> This household-income targeted subsample is constructed due to the requirements of an unrelated study with questions that are fielded on this survey.

any treatment effects we observe.<sup>11</sup> Thus, the estimated average treatment effects we report for the Lucid sample are estimates of the population average treatment effect ( $\widehat{PATE}$ ). We acknowledge, however, that it is possible that our sampling weights could be misspecified—e.g., there may be factors other than those that are accounted for when constructing the weights that systematically predict selection into the sample. As such, we note any meaningful differences between weighted and unweighted results ( $\widehat{PATE}$  versus sample average treatment effect  $\widehat{SATE}$ ) and show all unweighted analyses in the Online Appendix.<sup>12</sup> Furthermore, we categorize the outcomes of interest into families by the underlying concern that is being measured, and use the Benjamini–Hochberg correction ( $\alpha = 0.05$ ) for multiple comparisons within each family (Benjamini and Hochberg 1995). Details regarding all the procedures we follow in our analyses were decided and documented prior to data collection in a pre-analysis plan.<sup>13</sup>

Survey respondents are asked about two hypothetical prospective policies: a program that provides newborns and their mothers with nutritional assistance (*nutrition policy*) and a program that helps employed workers get additional training to remain competitive in the job market (*job training policy*). The question wording and design are shown in Table 1. The descriptions of both policies resemble those that currently exist—the nutrition policy is similar to the Special Supplemental Nutrition Program for Women, Infants, and Children (WIC) and the job training policy is similar to many programs that are funded through the US Department of Labor.<sup>14</sup> However, there are important differences between these hypothetical policies and their real-world analogs. The proposed nutrition policy assists *all* newborns and their mothers, and therefore unlike WIC, benefits are not targeted toward any particular income-group. Correspondingly, the proposed job training policy allows those who are *employed* to choose how best to supplement their skills whereas existing policy largely targets those who have lost their jobs. Thus, each policy presented to respondents is more universalistic than existing programs and therefore designed

---

<sup>11</sup> We use a raking procedure to match the marginal distributions of our sample to the ACS margins of age, education, household income, gender, and minority identification. See, Battaglia et al. (2009) for an overview of raking as well as practical considerations for implementing the procedure.

<sup>12</sup> For a discussion on estimating and interpreting survey experimental treatment effects with survey weights, see Miratrix et al. (2018) and Franco et al. (2017).

<sup>13</sup> Analyses of the Lucid sample and SSI sample (discussed in Experiment 2) reflect decisions made prior to data collection and that were documented in two pre-analysis plans registered with Evidence in Governance and Politics (EGAP). The pre-analysis plans are filed jointly under EGAP ID 20170725AA. Materials to replicate the analyses presented in this article are available here: <https://doi.org/10.7910/DVN/XYAQC>.

<sup>14</sup> WIC is a federally funded program that targets low-income pregnant and postpartum women, infants, and children up to 5 years of age. Depending on the State agency administering the benefit, participants receive WIC benefits through checks, vouchers, or as electronic benefit transfers (see, <https://www.fns.usda.gov/wic/>). A typical example of an existing US Department of Labor job training policy is the Trade Adjustment Assistance Program (TAA). The TAA targets workers that have lost their jobs or have their jobs threatened due to trade-related circumstances. Eligible individuals receive some combination of classroom training, on-the-job training, income support, and reimbursement for job search related expenses (see, [https://www.doleta.gov/tradeact/docs/program\\_brochure2014.pdf](https://www.doleta.gov/tradeact/docs/program_brochure2014.pdf)).



not to prime any views regarding the economic standing or deservingness of potential beneficiaries.

The manipulated feature of each hypothetical policy stipulates how benefits will be distributed to citizens. In the case of the nutrition policy, \$100 each month will be transferred to households by either a debit card or as tax reimbursements so long as the funds are used to buy approved items. We specify the debit card as the direct spending treatment as it's most analogous to a cash transfer.<sup>15</sup> In contrast, the indirect treatment of the nutrition policy mirrors the modal “submerged” government programs discussed previously—households receive a tax credit. We specify a tax credit as opposed to a tax deduction to ensure that households with no tax liability or insufficient liquidity are still able to receive benefits. Turning to the job training policy, we use a similar logic as in the nutrition policy treatments. Workers receive benefits through a bank transfer or as a tax credit.<sup>16</sup>

Immediately following the complete description of the policy including the randomly assigned treatment, respondents are asked their level of support for the program. We also probe the underlying considerations that could drive any differential support by type of transfer. To this end, we ask respondents about the costs associated with the program and their perceptions of prospective beneficiaries. In terms of costs, respondents estimate the percentage of eligible individuals who would use the policy as well as the number of cents out of each dollar spent on that policy that would be diverted away for purposes other than the policy goal. We also explicitly ask respondents to compare the type of transfer they receive in the treatment to the alternative—those assigned to see the direct treatment are asked whether this mode of delivery would cost the government more than an equivalent indirect transfer, and vice versa.<sup>17</sup> To measure perceptions of prospective beneficiaries, we ask respondents if particular income-groups would benefit more than others under the policy, the extent to which households who (do not) claim benefits actually (do not) need benefits, whether beneficiaries need to expend effort to receive benefits, and if the program would lead to a greater dependency on government.

Panels a and b of Fig. 1 show average levels of support for the nutrition and job training policies under each spending treatment condition, respectively. Comparing levels of support between treatments, we observe that respondents are substantially more supportive of indirect spending over direct spending for both the nutrition and job training policies (Nutrition: Difference-of-means,  $\widehat{PATE} = 0.125$  in increased support,  $p = 0.067$ ; Job Training:  $\widehat{PATE} = 0.136$  in increased support,  $p < 0.05$ ). The increase in support when the government transfers benefits indirectly is approximately 10% and 12% of the standard deviation of responses in the direct spending

<sup>15</sup> Note, that this method of transfer is similar to Electronic Benefit Transfer (EBT) cards currently used by many states.

<sup>16</sup> Note, however, that certain benefits under TAA are transferred to eligible individuals as a reimbursement for a qualified expense.

<sup>17</sup> Given that exposing respondents to both direct and indirect framings of the policy could generate biases in subsequent responses, we present this question last, and only in regards to the final policy that is shown. E.g., a respondent assigned to view the nutrition policy followed by the job training policy is only asked this question about the job training policy at the end of the corresponding question-block.

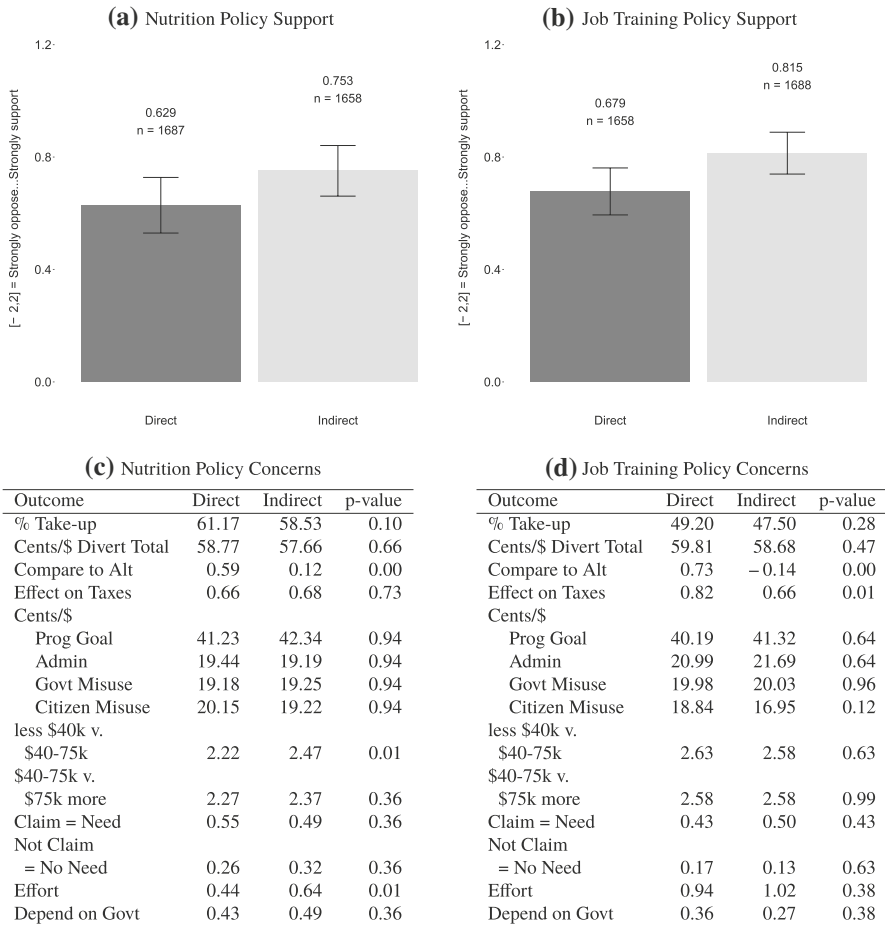


**Table 1** Experiment 1 question wording and design

Policy-treatment	Vignette wording
Nutrition policy description	Would you support a new government program that provides all newborns and their mothers with nutritional assistance for the first year after birth?
AND	
Direct treatment	Households with newborns would receive \$100 (untaxed) on a debit card each month that could be used to buy approved nutritional items.
OR	
Indirect treatment	Households with newborns could receive a tax credit of \$100 each month, as long as they spend that much on approved nutritional items.
Job training policy description	Would you support a new government program that helps workers who currently have a job to get additional training or a college degree in order to improve their productivity and remain competitive in the job market should they be unemployed in the future?
AND	
Direct treatment	Employed individuals would receive \$7000 (untaxed) deposited in their checking account each year, for up to 4 years, as long as the money is spent on approved training programs or college classes.
OR	
Indirect treatment	Employed individuals could receive a tax credit of \$7000 each year, for up to 4 years, as long as they spend that much on approved training programs or college classes.
Complete vignettes consist of the policy description followed by either the direct or indirect treatment. Respondents see both policies and are randomly assigned to one of two treatment arms:	
(1) Direct Nutrition Policy and Indirect Job Training Policy	
(2) Indirect Nutrition Policy and Direct Job Training Policy	
Complete vignettes are presented in random order.	

treatment groups for the nutrition and job training policies, respectively. Thus, respondents appear to hold a similar preference for indirect spending across the two policy domains.

The type of transfer may affect beliefs about the costs associated with the program, a potential source of variation in public support. We ask respondents to consider the extent to which people will actually use the policies at hand. Respondents believe that indirect transfers would lead to less take-up than direct transfers (approx. 2% points for each policy, though the effect is not statistically significant ( $\alpha = 0.05$ )). This belief regarding use might reflect underlying perceptions about the added costs to the government. Indeed, we find that when respondents are asked to compare the two modes of delivery to each other, they overwhelmingly view



**Fig. 1** Experiment 1 results. **a, b** show levels of support by treatment measured on a 5-point scale: - 2 = Strongly oppose...0 = Neither support nor oppose...2 = Strongly support. Panels *c* and *d* show other outcomes by treatment: “% Take-up” are avg. perc. of eligible beneficiaries that would use the policy; “Cents/\$Diverted Total” are avg. perc. of govt. spending on the policy that would be diverted away from the program goal; “Compare to Alt” are avg. responses to whether the treatment program would cost govt. more than the alternative (- 2 = program costs lot less than alternative...2 = program costs lot more than alternative); “Effect on Taxes” are avg. responses to the perceived effect of the program on taxes (- 2 = lower my taxes a lot...2 = raise my taxes a lot); “Cents/\$Prog Goal” are avg. perc. of govt. spending on the policy that would be spent on the program goal; “Cents/\$[Admin, Govt Misuse, Citizen Misuse]” are avg. perc. of govt. spending on the policy that would be spent on govt. admin. (bureaucracy) and inefficiency, govt. misuse and abuse, and citizen (and business for the job training policy) misuse and abuse; “less \$40k vs. \$40-75k” and “\$40-75k vs. \$75k more” report avg. responses to the question of which income-group within each pair benefits much more under the policy (1 = 1st group benefits much more...5 = 2nd group benefits much more); “Claim = Need” and “Not Claim = No Need” are avg. responses to whether households/workers (not) claiming benefits do (not) really need them, “Effort” are avg. responses to whether receiving benefits requires households/workers to put in some effort, “Depend on Govt” are avg. responses to whether the program will lead to households/workers expecting govt. help in similar circumstances (- 2 = Str. disagree...2 = Str. agree). **a, b** show 95% bootstrapped confidence intervals; **c, d** report two-sided bootstrapped p-values, adjusted using the Benjamini–Hochberg correction for multiple comparisons. Details regarding the procedure we use were decided and documented prior to data collection in a pre-analysis plan

indirect transfers as less costly to the government, relative to direct transfers (Nutrition: Difference-of-means,  $\widehat{PATE} = -0.47$  in decreased cost,  $p < 0.01$ ; Job Training: Difference-of-means,  $\widehat{PATE} = -0.87$  in decreased cost,  $p < 0.01$ ; see row 3 of Fig. 1c, d). And, furthermore, while respondents believe that both programs will lead to an increase in their taxes, in the case of the job training policy, indirect delivery is associated with a smaller increase in taxes relative to direct delivery (Difference-of-means,  $\widehat{PATE} = -0.16$  in decreased taxes,  $p = 0.01$ ; see row 4 of Fig. 1d).

Perceptions of cost could also manifest as concerns regarding the potential for waste, fraud, and abuse when contemplating prospective government spending. We therefore ask respondents to estimate the number of cents out of each dollar that would be diverted away from the program goal. Specifically, respondents are asked to divide a dollar of spending into the following line-items: providing nutritional assistance to newborns/mothers *or* helping workers get additional training/college degree, government administration (bureaucracy) and inefficiency, government misuse and abuse, and citizen/[business] misuse and abuse.<sup>18</sup> Rows 5–8 of panels c and d of Fig. 1 show these results. It is worth noting that respondents report that close to 60% of government revenues will be diverted away from the actual benefit in both policies. However, in terms of the total percentage diverted as well as for the individual line-items, there are no substantively or statistically significant differences by treatment.

Apart from considerations of costs, respondents' support could also reflect perceptions about the intended beneficiaries under each policy. Generically, respondents believe that both policies are slightly skewed towards lower-income households. However, in the case of the nutrition policy, we find that respondents view middle-income households as benefiting more than the lowest-income households when the transfer is indirect (Difference-of-means,  $\widehat{PATE} = 0.24$  in increased benefits,  $p = 0.01$ ; see row 9 of Fig. 1c). We do not find an equivalent perception with respect to the job training benefit. Similarly, when thinking about the effort citizens must expend to receive program benefits, respondents believe that the indirect delivery of the nutrition policy requires greater effort by beneficiaries relative to the direct transfer (Difference-of-means,  $\widehat{PATE} = 0.20$  in increased effort,  $p = 0.01$ ; see row 13 of Fig. 1c)—no such perception exists in regards to job training. Additionally, in terms of citizen need, respondents do not perceive one type of delivery as more likely to ensure that citizens in need will actually claim benefits from the program. Nor do respondents believe that program delivery could differentially affect citizens' expectations that government will help in similar circumstances.

In summary, citizens clearly view direct and indirect spending differently when thinking about prospective programs. While respondents strongly prefer that the government deliver benefits indirectly, on average (consistent with Haselswerdt and Bartels 2015), we identify potential reasons for this opinion that pertain to cost and views regarding potential beneficiaries. First, respondents believe that indirect transfers will cost the government less than direct transfers, on average. In the case of the job training policy, respondents may link this overarching view about cost to their

<sup>18</sup> For the job training policy respondents are asked about "Citizen/business misuse and abuse" so as to include any funds diverted by educational institutions.

personal tax liability—indirect delivery is viewed as increasing taxes less compared to the direct version of the program. Second, when considering the nutrition policy, respondents view indirect delivery as being more beneficial to middle-income households over those with the lowest incomes. And, respondents believe that beneficiaries of the nutrition policy must put in greater effort when the benefit is delivered indirectly versus directly.

## Experiment 2: Evidence that Regressivity Does Not Affect Preferences

While citizens favor indirect government spending when contemplating certain new policies, what are their preferences regarding existing programs? Mettler (2011) and Faricy and Ellis (2013) argue that if the distributional consequences of indirect programs were known to the public—that is, made visible—those that are regressive would become less popular. In a second survey experiment, we test this proposition by providing respondents information about the beneficiaries of the Home Mortgage Interest Deduction (HMID) and the Unemployment Insurance (UI) system before measuring support for these policies. We choose the HMID as it is a large *indirect* benefit delivered as a tax expenditure and a program that has been the subject of similar research designs. Conversely, the UI system provides a *direct* benefit and has not received as much scholarly attention as the HMID.

Using the same sample from Experiment 1, we employ an experimental design that manipulates the distributional consequences of a policy. Respondents are randomly assigned to see either information about the HMID or the UI system. The vignettes and design are shown in Table 2. Those assigned to view information about the HMID are presented with a short description of the policy followed by a regressive *or* progressive portrayal of the distributional effects of the policy. The regressive treatment states that eligible households earning between \$75,000 and \$200,000 a year receive approximately twice as much in benefits, in dollars, as households earning less than \$75,000. In the progressive treatment, the benefit is expressed as a percentage of the average incomes earned by each group. Thus, households earning less than \$75,000 a year receive twice as much in benefits relative to their income as compared to households earning between \$75,000 and \$200,000.<sup>19</sup>

Respondents assigned to view information about the UI system are similarly presented with a short description of the policy followed by either a regressive or progressive treatment. We use the fact that UI premiums are capped at a lower annual income than UI benefits to frame the program as being regressive, while exploiting the cap on benefit amounts to frame the program as progressive.<sup>20</sup> Those assigned to the regressive treatment are told that eligible unemployed workers who had earned

<sup>19</sup> HMID benefit figures are based on authors' calculations using the following data: Estimates of Federal Tax Expenditures for Fiscal Years 2015–2019 (Joint Committee on Taxation), Table 1.2. All Returns: Adjusted Gross Income, Exemptions, Deductions, and Tax Items, by Size of Adjusted Gross Income and by Marital Status, Tax Year 2014 (IRS). Both figures are conditional on take-up, but higher income households are much more likely to take up the policy.

<sup>20</sup> UI benefit figures are based on authors' calculations using US Department of Labor data: Significant Provisions of State Unemployment Insurance Laws, Effective January 2014.

more than \$250,000 a year get more benefits per dollar of the premiums they pay compared to those who had earned less than \$25,000. In the progressive treatment, the benefit is again expressed as a percentage of the average incomes earned by each group. Thus, workers who had earned less than \$25,000 a year receive a substantially larger benefit relative to their pre-job loss income than workers who had earned more than \$250,000.

As we discuss in detail above, both Mettler (2011, Ch. 3) and Faricy and Ellis (2013) examine support for the HMID. Our experimental design has a critical difference from these previous studies that allow us to make more precise inferences regarding the effect of distributional information on citizen preferences. In particular, our design allows us to compare outcomes for respondents who receive regressive versus progressive framings of the same policy. This is a key difference from the design used in the previous studies—both prior studies compare responses between a treatment where no distributional information is shown and a treatment where a regressive framing is shown. Thus, our experiment isolates the effect of the particular type of distributional framing—regressive or progressive—without conflating it with the generic effect of priming any distributional considerations.

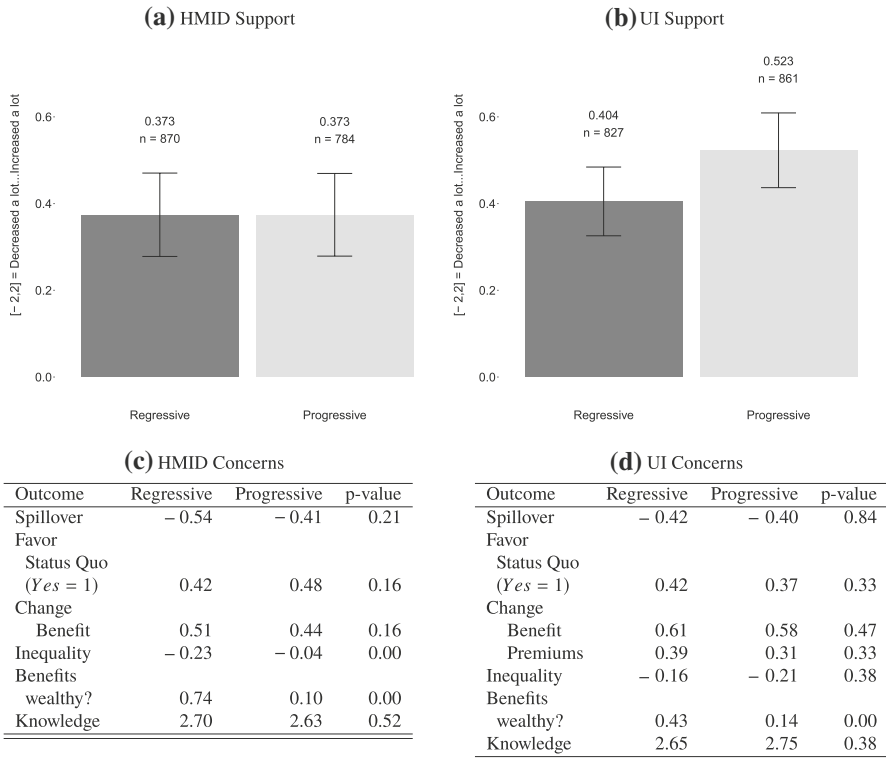
After respondents are shown a complete vignette describing the HMID or UI system, we first ask whether benefits provided under the program should be increased, decreased, or kept about the same. We also ask an analogous question with respect to government activity in general. Respondents may also want to alter the distributive features of the program rather than favoring expansion or reduction. To this end, we next ask respondents whether they favor changing, abolishing, or keeping the program about the same. In the HMID policy arm, respondents can also state a preference to change benefit amounts for low and high-income households, while keeping the overall cost of the program fixed. In the UI policy arm, respondents are able to change both the benefit amounts and the premium contributions that individuals of different incomes receive and pay, respectively.

Panels a and b of Fig. 2 show the levels of support by the distributional treatments. We find no meaningful difference, on average, in respondent support between regressive and progressive framings of the HMID. However, in the case of the UI system, respondents are more supportive of the program when it is framed progressively (Difference-of-means,  $PATE = 0.12$  in increased support,  $p < 0.05$ ). It is worth highlighting that this support response measures whether respondents want to increase or decrease spending on the program and is coded on a 5-point scale, between  $-2$  and  $2$ . Positive values therefore indicate a favorable view of a given policy. In both framings of the HMID and UI system, respondents want increased spending on the program. Even when framed regressively, these policies are popular. We also find that the framings have no meaningful effect on respondents' demands for changing the level of government activity in general (see row 1 of Fig. 2c, d, labeled "spillover").

We make sure that respondents understood the distributional information presented to them by asking two construal questions. In both policy arms, respondents who view the regressive frame believe that the policy benefits wealthy citizens more than low-income citizens, on average, compared to respondents who view the progressive frame (see row 5 of Fig. 2c and row 6 of Fig. 2d). And, in the case of the HMID, respondents assigned to the regressive frame believe that the policy

**Table 2** Experiment 2 vignette wording and design

Policy-treatment	Vignette wording
HMID policy description	The home mortgage interest deduction is a current government program that helps individuals and their households/families buy their own homes. A household that is paying a mortgage on their home receives benefits from this program. This benefit is delivered by allowing the household to deduct the amount they pay in interest on their mortgage from their adjusted gross income on their federal tax forms.
AND	
Regressive treatment	While all homeowners can qualify for the program, the benefit for eligible households earning between \$75,000 and \$200,000 a year is, on average, about twice as large (in dollars) as the benefit for eligible households earning less than \$75,000.
OR	
Progressive treatment	While all homeowners can qualify for the program, the benefit for eligible households earning less than \$75,000 a year is, on average, about twice as large (as a percentage of the incomes they earn) as the benefit for eligible households earning between \$75,000 and \$200,000.
UI policy description	The Unemployment Insurance system is a current government program that helps people who have lost their jobs by temporarily replacing part of their wages while they look for work. This program is funded by a tax workers pay on their wages. On average, individuals who earn more than \$17,000 a year pay a fixed \$476 each year as an insurance premium when they are employed. An individual who loses their job receives a benefit from this program for up to 26 weeks.
AND	
Regressive treatment	However the benefit for eligible individuals who earn more than \$250,000 is, on average, substantially larger (as a percentage of the premiums they pay) than the benefit for eligible individuals who earn less than \$25,000.
OR	
Progressive treatment	However, the benefit for eligible individuals who earn less than \$25,000 is, on average, substantially larger (as a percentage of the incomes they earn) than the benefit for eligible individuals who earn more than \$250,000.
Complete vignettes consist of the policy description followed by a regressive/progressive treatment. Respondents are randomly assigned to see one policy.	



**Fig. 2** Experiment 2 distributive treatment results. **a, b** show levels of support (spending on program) measured on a 5-point scale: - 2 = Decreased a lot, - 1 = Decreased, 0 = Kept about the same, 1 = Increased, 2 = Increased a lot. **c, d** show other related outcomes by treatment: “Spillover” are avg. responses to whether the government activity should be increased/decreased (- 2 = govt. activity should be decreased a lot and you may end up paying less in taxes...2 = govt. activity should be increased a lot and you may end up paying more in taxes); “Favor Status Quo” are the proportions of respondents that favor keeping the policy about the same; “Change Benefit” are the avg. responses to whether lower-income households/individuals should get more benefits: lower-income households/individuals should get more benefits (= 1), benefits should be kept about the same (= 0), lower-income households/individuals should get less benefits (= -1); “Change Premiums” are the avg. responses to whether lower-income individuals should pay a lower premium: lower-income individuals should pay a lower premium (= 1), premiums should be kept about the same (= 0), lower-income individuals should pay higher premium (= -1); “Inequality” are avg. responses to whether the policy increases econ. ineq. (= -1), keeps econ. ineq. about the same (= 0), decreases econ. ineq. (= 1); “Benefits wealthy?” are avg. responses to whether the program benefits wealthy individuals more than low-income individuals (- 2 = Strongly disagree...2 = Strongly agree); “Knowledge” are avg. responses to how much the respondent knows about the program: Nothing at all (1), a little (2), a moderate amount (3), a lot (4), a great deal (5). **a, b** show 95% bootstrapped confidence intervals; **c, d** report two-sided bootstrapped p-values, adjusted using the Benjamini–Hochberg correction for multiple comparisons. Details regarding the procedure we use were decided and documented prior to data collection in a pre-analysis plan

increases inequality more than those assigned to the progressive frame (see row 4 of Fig. 2c)—we do not find an effect in the UI system policy arm. Thus, despite concerns that the HMID benefits the wealthy and increases inequality when the policy is framed regressively, respondents’ support remains unchanged.



After learning about the distributional consequences of both policies, respondents may want to modify how much different income-groups benefit. We find suggestive evidence of this in the case of the HMID. Respondents assigned to the regressive treatment are slightly less likely to favor keeping the policy as is (see row 2 of Fig. 2c,  $p = 0.16$ ), on average, compared to respondents assigned to the progressive treatment. Furthermore, respondents assigned to the regressive treatment are slightly more supportive, on average, of changing HMID benefits such that lower-income households receive more and higher-income households receive less compared to respondents assigned to the progressive treatment (see row 3 of Fig. 2c,  $p = 0.16$ ).<sup>21</sup> These results are not statistically significant at conventional levels and are not replicated in the UI policy arm.<sup>22</sup>

We show that regressive framings of the HMID and UI system do not, on average, cause respondents to support decreasing government spending on these programs. However, in the case of the UI system, respondents are substantially more supportive of the policy, on average, when it is framed progressively, relative to when the policy is framed regressively. While respondents in both policy arms construe the treatments as they were intended, this only resulted in a modest and insignificant increase in respondents' willingness to change the distribution of benefits for the HMID.

### Does Direct Spending Affect Support for Current Policies?

In Experiment 1 we show that citizens have a clear preference for the indirect delivery of prospective government benefits. Is this also the case for existing policies? Mettler (2011, Ch. 2) notes that a large percentage of Americans who say they oppose government spending are themselves beneficiaries of social programs, arguing that this may be the result of submerged policies that “mask” the role of government. An implication of this claim is that the public may view current government spending more favorably if benefits were transferred to citizens in a more direct manner—Faricy and Ellis (2013) attempt to test this for the HMID but are underpowered to draw a definitive conclusion. We therefore design and implement a second manipulation to our second study to see if a direct framing affects support for the HMID.

After viewing a description of the HMID (see row 1 of Table 2), respondents are randomly assigned to see either a direct or indirect framing of the policy. In the direct treatment, the households are described as receiving a cash transfer, equivalent to the deduction they claim:

<sup>21</sup> We find that treatment effect on preferences to modify the HMID are statistically significant ( $\alpha = 0.05$ ) when estimate the *SATE* without using weights. See, rows 3–4 of Online Appendix Figure B.1c.

<sup>22</sup> We do find, however, that the unweighted effect of the regressive frame is to increase support for lower-income individuals paying a lower UI premium is ( $\overline{SATE} = 0.08$ ,  $p < 0.05$ ). See, row 5 of Online Appendix Figure B.1d.

...It is equivalent, however, to the federal government depositing that benefit amount (untaxed) into their checking account each year as a subsidy for their mortgage. ...

Conversely, in the indirect treatment, households are described as receiving a reduction in taxes (the status quo):

...It is equivalent, however, to the federal government reducing the taxes they pay by that benefit amount each year as a subsidy for their mortgage. ...

This design is fielded on two different samples. The first sample consists of respondents drawn from a nationally representative online panel constructed by Survey Sampling International (SSI). For the second sample, we recruited participants on Amazon's Mechanical Turk (MTurk).<sup>23</sup>

While the direct/indirect treatments were worded exactly the same across samples, the distributional information shown to the MTurk sample are different than those shown to SSI respondents. Respondents in the MTurk sample are randomly assigned to see one of four distributional treatments: a control condition (no information), a regressive condition showing that the average dollar amount of the HMID benefit increases with annual household income, a progressive condition showing that the average HMID benefit as a percentage of average income taxes paid decreases with annual household income, and a treatment that shows both progressive and regressive framings.<sup>24</sup> Alternatively, respondents in the SSI sample view only one of two distributional treatments: the same regressive condition as in the MTurk sample, and a progressive condition showing the average HMID benefit as a percentage of the *incomes earned* by households.<sup>25</sup> After viewing a complete treatment—both directness and distributional information manipulations—respondents are asked whether benefits provided through the HMID should be increased, decreased, or kept about the same.

To analyze the MTurk sample responses, we regress the response on the direct/indirect treatment indicator and indicators that correspond to the distributional treatments.<sup>26</sup> The coefficients on the indirect treatment indicator are reported in Table 3. We find somewhat different results in the different samples. In the SSI sample, respondents state a small but statistically insignificant preference for the indirect framing of the HMID relative to the direct treatment ( $\widehat{SATE} = 0.04$  in increased support,  $p = 0.15$ ). However, in contrast, respondents in the MTurk sample assigned to the indirect treatment are substantially more supportive of the policy as compared to respondents assigned to the direct treatment ( $\widehat{SATE} = 0.09$  in increased support,

<sup>23</sup> For a comparison between MTurk, student convenience, and ANES samples, see Berinsky et al. (2012).

<sup>24</sup> The complete wordings of the distributional treatments employed in the MTurk sample are shown in Online Appendix Figure B.8.

<sup>25</sup> The complete wordings of the distributional treatments employed in the SSI sample are shown in Online Appendix Table B.5.

<sup>26</sup> The experiment conducted on MTurk was intended as a pilot study, thus we did not pre-register a plan for analysis.

$p < 0.05$ ). This effect in the MTurk sample is approximately 15% of the standard deviation of responses in the direct treatment. While the respondents in both samples appear to prefer the indirect framing of the HMID, the difference in the magnitude of results leads us to be cautious in characterizing citizens preferences for the delivery of currently indirect policies.<sup>27</sup> Either way, however, we find at worst no opposition to indirect delivery and perhaps some evidence of a preference for it. This suggests that citizens' support for current government spending may not increase simply by changing the way that policy benefits are delivered.<sup>28</sup>

## Discussion

Prior work argues that citizen support for government policies is affected by the way policy benefits are transferred to the public and by the distributional consequences of the policy. Though researchers have found that citizens prefer that policy benefits be delivered indirectly—via the tax code—the causes and consequences of this preference have been largely unexamined. Furthermore, while scholars have been quick to note that learning about the regressive consequences of certain indirect programs diminishes citizens' support for them, the research designs employed in these studies are ill-equipped to support such a dispositive conclusion.

In this article, we build upon the current literature in two ways. First, we elucidate citizen preferences for indirect government spending by measuring support for two prospective policies—concerning childhood nutrition and job training—randomly varying the way in which the government transfers benefits to eligible individuals. Our results strengthen and clarify the findings of Haselswerdt and Bartels (2015). We find that citizens state a strong *primitive* preference for the indirect delivery of both policies. And, importantly, we investigate several underlying reasons for this preference for indirect delivery.

To summarize, we find the following with respect to potential mechanisms: Citizens view indirect transfers as substantially less costly to the government, relative to direct transfers. And, with respect to the job training policy, direct delivery is seen as increasing personal taxes more than if the program were delivered indirectly. However, there is no difference in citizens' perceptions of efficiency—funds diverted to waste, fraud, and abuse—when evaluating these direct and indirect programs. When considering potential beneficiaries of these prospective policies, for the nutrition

---

<sup>27</sup> The differences in results between samples could be the result of a number of factors. We call attention to two such factors. First, the type of respondent in each sample: MTurk respondents complete the survey as a “Human Intelligence Tasks (HITs)” in an online labor market whereas the average SSI respondent is usually compensated to take part in marketing research. Second, as the distributional treatments used in the MTurk sample are much different than those used in the SSI sample, the information about HMID beneficiaries presented to MTurk respondents could interact with the direct/indirect treatment in ways that are different than we observe with SSI respondents.

<sup>28</sup> We note that even when restricting our attention to the regressive framing of the HMID, respondents in both the direct and indirect treatment conditions favor, on average, increasing spending on the program.

**Table 3** Effect of direct versus indirect framing of HMID on program support

	$\widehat{SATE}_{SSI}$	<i>p</i>	$\widehat{SATE}_{MTurk}$	<i>p</i>
– 1 = Dec. a lot... 1 = Inc. a lot	0.04	0.15	0.09	0.02
n	1241		1066	

Table shows average differences between direct and indirect treatments (*SATE*). The response in the SSI sample is measured on a 5-point scale: – 1 = Decreased a lot, – 0.5 = Decreased, 0 = Kept about the same, 0.5 = Increased, 1 = Increased a lot. In the MTurk sample, this response is measured on a 3-point scale: – 1 = Decreased, 0 = Kept about the same, 1 = Increased. P values are computed from standard OLS models with appropriate treatment indicators and heteroskedasticity robust standard errors

program when benefits are indirectly delivered (rather than directly), citizens believe that middle-income households will be advantaged more than the lowest-income households. Additionally, indirect nutrition benefits are seen as requiring more effort on the part of households to receive the transfer, relative to direct benefits.

We note that real-world policy design is the result of political competition and legislative bargains. As such, citizens often learn about policy delivery through partisan messages. Furthermore, citizens’ partisan identity itself could be associated with varied motivations in preferring one type of policy design over another.<sup>29</sup> In order to examine this potentially important source of underlying variation in preferences for policy design, we note differences in treatment effects between citizens who identify as (or lean towards) Democrats or Republicans. In the case of the nutrition policy, we find that Republicans have a much stronger preference for indirect delivery than do Democrats. However, both groups share a similar belief that indirect delivery costs the government less, privileges middle-income households more than lowest-income households, and requires more effort on the part of beneficiaries, compared to a direct transfer (see Online Appendix Figure A.1).

When considering the job training policy, however, Democrats are essentially indifferent between direct and indirect program delivery—Republicans, on the other hand, express a strong preference for an indirect transfer. Both groups express a similar belief that the cost of the job training program will be less when benefits are transferred indirectly to workers. And, Republicans further associate indirect delivery with a smaller increase in their own taxes, lower amounts of citizen misuse, and a lower likelihood that citizens will become dependent on government, compared to a direct transfer (see Online Appendix Figure A.2). Overall, neither partisan group supports direct over indirect delivery and Republicans state a strong preference against direct transfers.

Our second contribution is to test whether providing citizens information regarding the distributive consequences of existing policies—utilizing both regressive

<sup>29</sup> For discussions on the effects of partisanship on different apolitical activities, see, Margolis and Sances (2016), Gerber and Huber (2009), Oliver et al. (2015), among others.

and progressive frames—alters support for them. Contrary to Mettler (2011, Ch. 3) and Faricy and Ellis (2013) we find little evidence that framing the HMID as having regressive distributional effects generates a big shift in support for the policy in comparison to the progressive framing. We do, however, find that citizens are more supportive of the UI system when it is progressively framed. But, in considering both the HMID and UI system, citizens favor increasing benefits under both distributional frames. We also find suggestive evidence that when citizens learn about the regressive distributional consequences of the HMID, they are less likely to favor the status quo program design and are more likely to favor making benefits more progressive. Indeed, we note that overall, Democrats state a strong preference to make both the HMID and UI system more progressive, whereas Republicans' preferences are largely unaffected by the distributional treatment (see Online Appendix Figures A.3, A.4).<sup>30</sup>

Our findings have several implications that advance the study of policy design and policy feedbacks. We show that citizens prefer indirect government spending, associating such transfers with lower overall costs to the government relative to direct transfers. However, we find no evidence that the type of transfer affects citizen perceptions of government revenues being diverted towards waste, fraud, and abuse. Thus, while citizens do appear to consider the budgetary implications of policy design, they do not seem to give much weight to efficiency concerns. This is indeed puzzling given recent research that privileges citizen distrust of government as a leading factor for the decline in public support for social spending (Hacker and Pierson 2017; Kuziemko et al. 2015). At the very least, such distrust does not appear to be differentially activated by how program benefits are delivered.

Additionally, our finding that distributive information does little to affect support for current programs highlights an important limitation for real-world political rhetoric. While elites on the left may try to mobilize public support against regressive government programs, our evidence shows that these efforts may not be successful in part because these programs are popular even when regressive. We do find that proposals that make regressive policies more progressive, in terms of both the distribution of costs and benefits, could gain substantial traction with the electorate, but the key may be to emphasize marginal changes rather than wholesale program removal.<sup>31</sup>

It is also worth noting that debates around the distribution of costs and benefits of government programs often take place in the context of political competition (e.g., elections). Thus, it is likely that regressive framings of current policies utilized by the left will be matched by countervailing progressive framings on the right. Our design and subsequent results highlight the importance of operationalizing both

<sup>30</sup> We do find, however, that the unweighted effect of the regressive frame is to increase support for lower-income individuals paying a lower UI premium among Republicans is ( $\widehat{SATE}_{Rep} = 0.12$ ,  $p < 0.05$ ). See, row 5 of Online Appendix Figure A.4d.

<sup>31</sup> A potential real-world example can be found in recent changes to the HMID that lowers the limit on qualifying mortgages from \$1,000,000 to \$750,000. See: [http://wapo.st/2zMo5QV?tid=ss\\_mail&utm\\_term=.e57003cdcdc7](http://wapo.st/2zMo5QV?tid=ss_mail&utm_term=.e57003cdcdc7).

such frames in order to understand the net effect of competing arguments on the policy preferences of citizens.

In summary, these findings extend previous survey and survey-experimental work by examining the effects of policy design on public opinion and by probing the specific considerations that may cause citizens to favor a particular means of policy delivery. We are aware of the limitations of a survey experimental approach as a substitute for the actual experience citizens have when experiencing policies delivered in different ways (i.e., as a substitute for taking a tax deduction or receiving a cash transfer). Nonetheless, we believe our account takes a first step in a promising new line of inquiry into the mass public's preferences and motivations that may shape the design of government spending programs.

**Acknowledgements** We thank Sarah Anzia, Christopher Berry, Natália S. Bueno, Alan Gerber, Jacob Hacker, John Henderson, Reuben Kline, Nolan McCarty, Patrick Tucker, Ebonya Washington, the three anonymous reviewers of this article and seminar participants at Yale University, the 2018 Emory University Conference on Institutions and Law Making, the 2018 New York University CESS Experimental Political Science Conference, and the 2017 American Political Science Association annual meeting for helpful comments and advice.

## References

- Applebaum, L. D. (2001). The influence of perceived deservingness on policy decisions regarding aid to the poor. *Political Psychology*, 22(3), 419–442.
- Arnold, R. D. (1992). *The logic of congressional action*. New Haven: Yale University Press.
- Battaglia, M. P., Hoaglin, D. C., & Frankel, M. R. (2009). Practical considerations in raking survey data. *Survey Practice*, 2(5), 1–10.
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 57(1), 289–300.
- Berinsky, A. J., Huber, G. A., & Lenz, G. S. (2012). Evaluating online labor markets for experimental research: Amazon.com's Mechanical Turk. *Political Analysis*, 20(03), 351–368.
- Campbell, A. L. (2003). *How policies make citizens: Senior political activism and the American welfare state. Princeton studies in American politics: Historical, international and comparative perspectives*. Princeton: Princeton University Press.
- Charness, G., Gneezy, U., & Kuhn, M. A. (2012). Experimental methods: Between-subject and within-subject design. *Journal of Economic Behavior and Organization*, 81(1), 1–8.
- Dunning, T. (2012). *Natural experiments in the social sciences: A design-based approach. Strategies for social inquiry*. Cambridge: Cambridge University Press.
- Ellis, C., & Stimson, J. A. (2012). *Ideology in America*. Cambridge: Cambridge University Press.
- Farcy, C., & Ellis, C. (2013). Public attitudes toward social spending in the United States: The differences between direct spending and tax expenditures. *Political Behavior*, 36(1), 53–76.
- Franco, A., Malhotra, N., Simonovits, G., & Zigerell, L. J. (2017). Developing standards for post-hoc weighting in population-based survey experiments. *Journal of Experimental Political Science*, 4(2), 161–172.
- Gerber, A. S., & Huber, G. A. (2009). Partisanship and economic behavior: Do partisan differences in economic forecasts predict real economic behavior? *American Political Science Review*, 103(03), 407–426.
- Gilens, M. (1999). *Why Americans hate welfare. Race, media, and the politics of antipoverty policy*. Chicago: University of Chicago Press.
- Gilens, M. (2009). Preference gaps and inequality in representation. *Political Science & Politics*, 42(02), 335–341.

- Hacker, J. S., & Pierson, P. (2010). *Winner-take-all politics. How Washington made the rich richer and turned its back on the middle class*. New York: Simon and Schuster.
- Hacker, J. S., & Pierson, P. (2017). *American amnesia. How the war on government led us to forget what made America prosper*. New York: Simon and Schuster.
- Haselswerdt, J., & Bartels, B. L. (2015). Public opinion, policy tools, and the status quo: Evidence from a survey experiment. *Political Research Quarterly*, 68(3), 607–621.
- Henry, P. J., Reyna, C., & Weiner, B. (2004). Hate welfare but help the poor: How the attributional content of stereotypes explains the paradox of reactions to the destitute in America. *Journal of Applied Social Psychology*, 34(1), 34–58.
- Hetherington, M. J. (2005). *Why trust matters. Declining political trust and the demise of American liberalism*. Princeton: Princeton University Press.
- Hetherington, M. J., & Rudolph, T. J. (2008). Priming, performance, and the dynamics of political trust. *The Journal of Politics*, 70(2), 498–512.
- Howard, C. (2007). *The welfare state nobody knows: Debunking myths about U.S. social policy*. Princeton: Princeton University Press.
- Jacoby, W. G. (1994). Public attitudes toward government spending. *American Journal of Political Science*, 38(2), 336–361.
- Kuziemko, I., Norton, M. I., Saez, E., & Stantcheva, S. (2015). How elastic are preferences for redistribution? Evidence from randomized survey experiments. *The American Economic Review*, 105(4), 1478–1508.
- Margolis, M. F., & Sances, M. W. (2016). Partisan differences in nonpartisan activity: The case of charitable giving. *Political Behavior*, 39(4), 839–864.
- McCarty, N., Poole, K. T., & Rosenthal, H. (2016). *Polarized America. The dance of ideology and unequal riches*. Cambridge: MIT Press.
- Mettler, S. (2011). *The submerged state: How invisible government policies undermine American democracy*. Chicago: University of Chicago Press.
- Miratrix, L. W., Sekhon, J. S., Theodoridis, A. G., & Campos, L. F. (2018). Worth weighting? How to think about and use weights in survey experiments. *Political Analysis*, 26(03), 275–291.
- Morgan, K. J., & Campbell, A. L. (2011). *The delegated welfare state: Medicare, markets, and the governance of social policy. Studies in postwar American political development*. Oxford: Oxford University Press.
- Oliver, J. E., Wood, T., & Bass, A. (2015). Liberellas versus conservatives: Social status, ideology, and birth names in the United States. *Political Behavior*, 38(1), 1–27.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.