

DO MEANS OF PROGRAM DELIVERY AND DISTRIBUTIONAL CONSEQUENCES AFFECT POLICY SUPPORT?

EXPERIMENTAL EVIDENCE ABOUT THE SOURCES OF CITIZENS' POLICY OPINIONS

Vivekinan L. Ashok*
Gregory A. Huber†

January 15, 2018

*Ph.D. Candidate, Yale University, Department of Political Science and Institution for Social and Policy Studies (vivekinan.ashok@yale.edu).

†Professor, Yale University, Department of Political Science; Resident Fellow, Institution for Social and Policy Studies (gregory.huber@yale.edu).

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 state a preference for indirect spending that is independent of the distributional consequences of a given policy. And, contrary to previous research, we find little evidence that highlighting the regressive effects of current government programs significantly diminishes their popularity. Our findings have implications for understanding the political calculus of policy design and the potential for public persuasion.

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.¹ 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.²

A potentially powerful alternative explanation for the preponderance of programs whose benefits are delivered indirectly is that citizens *prefer* indirect to direct delivery. Indirect delivery may be preferred because it is perceived as decreasing dependence on government, reducing opportunities for fraud or partisan corruption, less costly, or simply as a more efficient means of delivering

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

²Still 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. This effect of direct program delivery could be larger or smaller than the effect of directness on the understanding of the overall distribution of program costs and benefits.

benefits.

While some prior scholarship has sought to test both whether citizens prefer direct to indirect program delivery and how program support is affected by knowledge of the distribution of program costs and benefits, prior results are mixed. We present results from a new series of survey experiments that allow us 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 tentative explanations for this view. We also show that support for popular government spending, including the home mortgage interest deduction, is generally unaffected by whether the distributional consequences of the program are framed progressively or regressively, although information emphasizing program regressivity may cause citizens to support altering the distribution of program benefits to be more progressive.

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 exists despite evidence that citizens associate indirect spending with fewer individuals taking-up the programs when compared to equivalent direct spending. Furthermore, citizens do not perceive any difference in the aggregate level of fraud between transfer methods. In terms of perceived distributive effects of each policy, respondents do not believe that the mode of delivery of nutritional benefits advantages any particular income-group. However, in the case of job training, respondents perceive a strong regressive effect when transfers are made indirectly, but this policy design is still preferred to a direct transfer. Overall, these results suggest that elites may actually be reflecting constituent preferences when utilizing the tax code 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. Overall, we find little evidence that framing these policies as having regressive or progressive distributional effects generates a big shift in their popularity, contrary to Mettler (2011, Ch. 3) and Faricy and Ellis (2013). We do, find however, that citizens express a desire to make benefits and costs of these programs more progressive when they are framed regressively. Therefore, regressivity alone does not appear to endanger a program's survival, but does leave open the potential for attempts to make such programs more equitable.

PREVIOUS EXPERIMENTAL APPROACHES

Most scholarship that examines the political logic undergirding government spending design has relied on observational data and qualitative methods.³ 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 experiment 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

³See, e.g., Arnold (1992); Gilens (2009); Jacoby (1994); Howard (2007); Morgan and Campbell (2011); Ellis and Stimson (2012).

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.⁴ 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 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

⁴For a general discussion of the potential biases in within-subject designs, see, Charness, Gneezy and Kuhn (2012).

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.⁵

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. For example, it may be that individuals perceive there are different groups of beneficiaries depending on whether a program is administered as a direct cash transfer or via the tax code.

In light of these limitations, we designed and fielded three multi-arm survey experiments across two independent samples that examines support for both new and existing government programs while manipulating both the means of program delivery and the framing of the distributional

⁵For a general discussion of the challenges interpreting the effects of bundled treatments, see, Dunning (2012, pp. 300–302).

consequences of those programs (i.e., the same program is framed progressively or regressively).⁶ Additionally, we measure a variety of potential mechanisms examining how individuals' beliefs about, inter alia, program efficiency/waste, program take-up, and beliefs about group benefits, are affected by program design and the distributional frames.

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 Survey Sampling International (SSI) to direct a nationally representative online panel—by age, race, gender, region and partisanship—to a survey hosted on the Qualtrics survey platform.⁷ 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

⁶All analyses presented in this paper reflect decisions made prior to data collection that were documented in a pre-analysis plan.

⁷A number of recent articles in political science have used SSI samples. For examples, see: Berinsky, Margolis and Sances (2014); Malhotra, Margalit and Mo (2013); Malhotra and Margalit (2010). SSI recruits participants on an opt-in basis through various commercial websites. For more information about the construction of the panel, see: <https://www.surveysampling.com/solutions/data-collection/online-surveys/consumer/>.

Table 1: Experiment 1 Question Wording and Design

Policy - Treatment	Question 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?
<i>AND</i>	
Direct Treatment	Households with newborns would receive \$100 (untaxed) on a debit card each month that could be used to buy approved nutritional items.
<i>OR</i>	
Indirect Treatment	Households with newborns could file for reimbursements for up to \$100 each month as long as they spent that amount 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?
<i>AND</i>	
Direct Treatment	Employed individuals would receive \$7,000 (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.
<i>OR</i>	
Indirect Treatment	Employed individuals could reduce the taxes they pay by \$7,000 each year, for up to 4 years, as long as they spend that much on approved training programs or college classes.
<p><i>Complete questions consist of the policy description followed by either the direct or indirect treatment. Respondents see both policies, presented in random order, and are assigned to one of two treatment arms:</i></p> <p><i>(1) Direct Nutrition Policy and Indirect Job Training Policy</i></p> <p><i>(2) Indirect Nutrition Policy and Direct Job Training Policy</i></p>	

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.⁸ However, there are important differences between these hypothetical policies and their real-

⁸WIC 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

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 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 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.⁹ In contrast, the indirect treatment of the nutrition policy adds an element of friction between the benefit and the beneficiary—households must first purchase approved items and then file to be reimbursed. This *ex post* transfer is meant to change the order and nature of the interaction with the government, potentially obscuring the benefit relative to the direct treatment condition.¹⁰ Turning to the job training policy, workers receive the benefit through a bank transfer or via a tax deduction. We use a similar logic as in the nutrition policy treatments to specify the bank transfer as the direct spending treatment. However, the indirect transfer of job training benefits through a tax expenditure mirrors the modal “submerged” government program discussed above.

Immediately following the complete description of the policy including the randomly assigned treatment, respondents are asked their level of support for the program. Panels a and b of Figure 1 show average levels of support for nutrition and job training policies under each spending treatment condition, respectively. We also probe the underlying considerations that could drive any differential support by type of transfer. To this end, we ask respondents to estimate the percentage of eligible

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)).


⁹Note, that this method of transfer is similar to Electronic Benefit Transfer (EBT) cards currently used by many states.

¹⁰Note, however, that certain benefits under TAA are transferred to eligible individuals as a reimbursement for a qualified expense.

individuals who would use the policy as well as the number of cents out of each dollar spent on the policy that would be diverted for purposes other than the policy goal. Additionally, we ask respondents if particular income-groups would benefit more than others under the policy.¹¹

¹¹As discussed above, respondents are randomly assigned the order in which they are asked about each policy. We are able to replicate our findings if we restrict attention only to responses given to the first policy area shown.

Figure 1: Experiment 1 Results



Notes: Panels show average differences in outcomes between direct and indirect treatments for nutrition and job training policies. Panels *a* and *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 the avg. percentages of eligible beneficiaries that would use the policy; “Cents/\$Diverted Total” are the avg. percentages of government spending on the policy that would be diverted away from the program goal; “Cents/\$Prog Goal” are the avg. percentages of government spending on the policy that would be spent on the program goal; “Cents/\$[Admin, Govt Misuse, Citizen Misuse]” are the avg. percentages of government spending on the policy that are spent on government administration (bureaucracy) and inefficiency, government misuse and abuse, and citizen (and business for the job training policy) misuse and abuse; “less \$40k v. \$40-75k” and “\$40-75k v. \$75k more” report avg. responses to the question of which income-group within each pair benefits much more under the policy (1 = first group benefits much more, 5 = second group benefits much more). Panels *a* and *b* show 95% confidence intervals; panels *c* and *d* report p-values from two-sided, two-sample *t*-tests. P-values are 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.

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, $\widehat{ATE} = 0.139$ in increased support, $p < 0.01$; Job Training: $\widehat{ATE} = 0.163$ in increased support, $p < 0.01$). The increase in support when the government transfers benefits indirectly is approximately a quarter of the standard deviation of responses in the direct spending treatment groups for both policies. 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 proportion of eligible beneficiaries who end up using the program, a potential source of variation in public support. We therefore asked 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. 3% points for each policy, $p_{nutrition} = 0.03$, $p_{job_training} = 0.01$; see first row of Figures 1c, d). However, when we disaggregate potential beneficiaries by income-group, we find that respondents do not believe that the directness of the nutrition benefit would affect the type of household that would end up benefiting from the program. In the case of the job training benefit, by contrast, respondents view wealthier households as benefiting more than lower-income households when the transfer is indirect (see rows 7 & 8 of Figures 1c, d).

Perceptions of waste, fraud, and abuse might also shape preferences for a certain type of 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.¹² Rows 2–6 of panels c and d of Figure 1 show these results. It is worth noting that respondents believe 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

¹²For the job training policy respondents are asked about “Citizen/business misuse and abuse” so as to include any funds diverted by educational institutions.

no substantively or statistically significant differences by treatment.

In summary, citizens clearly view direct and indirect spending differently when thinking about prospective programs. While respondents strongly prefer that the government transfer revenues indirectly, on average (and consistent with Haselswerdt and Bartels (2015)), they appear to hold this opinion despite some concerns regarding the extent to which people will actually benefit from such transfers. First, respondents believe that fewer people will benefit from indirect spending, on average, compared to a direct approach. Second, in the case of the job training policy, indirect transfers are preferred even though the result may be more regressive than its direct equivalent.

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.

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.
<i>AND</i>	
Direct Treatment	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.
<i>OR</i>	
Indirect Treatment	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.
<i>AND</i>	
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.
<i>OR</i>	
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.
<i>AND</i>	
Regressive Treatment	However, the benefit for eligible individuals who earn more than \$75,000 is, on average, substantially larger (as a percentage of the premiums they pay) than the benefit for eligible individuals who earn between \$20,000 and \$40,000.
<i>OR</i>	
Progressive Treatment	However, the benefit for eligible individuals who earn between \$20,000 and \$40,000 is, on average, substantially larger (as a percentage of the incomes they earn) than the benefit for eligible individuals who earn more than \$75,000.
<p><i>Complete vignettes consist of the policy description followed by a direct/indirect treatment and a regressive/progressive treatment. Respondents are randomly assigned to see only one policy.</i></p>	

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 two experimental manipulations—a direct *or* indirect framing of the benefit, and 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 \$200,000.¹³ We reserve discussion of the direct/indirect treatment for the next section.

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.¹⁴ Those assigned to the regressive treatment are told that eligible unemployed workers who had earned more than \$75,000 a year get more benefits per dollar of the premiums they pay compared to those who had earned between \$20,000 and \$40,000. 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 between \$20,000 and \$40,000 a year receive a substantially larger benefit relative to their pre-job loss income than workers who had earned more than \$75,000.

As we discuss in detail above, both Mettler (2011, Ch. 3) and Faricy and Ellis (2013) examine

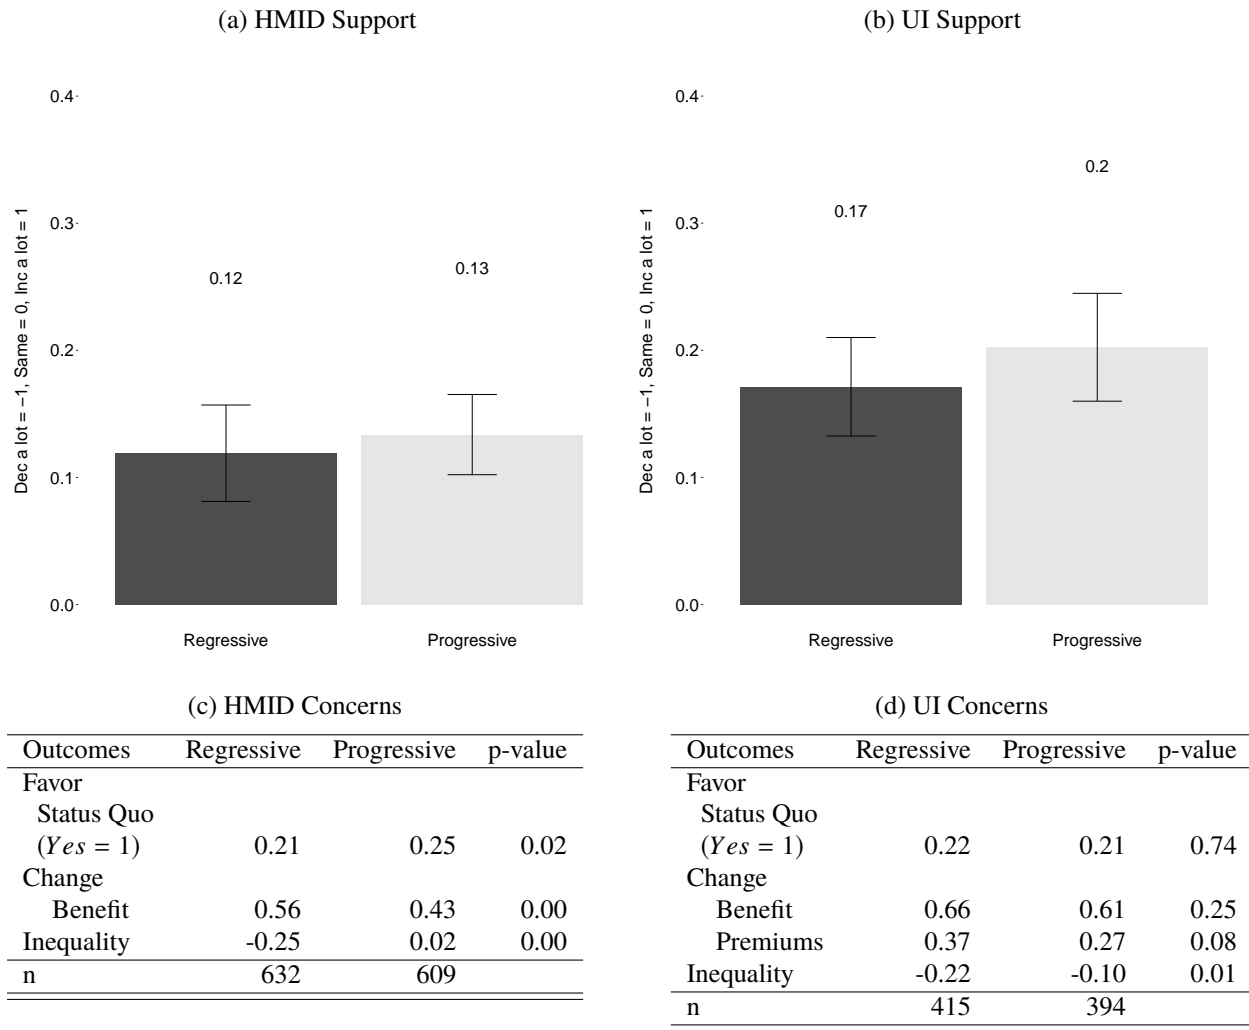
¹³HMID benefit figures are based on authors' calculations using the following data: Estimates of Federal Tax Expenditures for Fiscal Years 2015-2019 (Joint Committee on Taxation), Table 1.2. All Returns: Adjusted Gross Income, Exemptions, Deductions, and Tax Items, by Size of Adjusted Gross Income and 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.

¹⁴UI benefit figures are based on authors' calculations using US Department of Labor data: Significant Provisions of State Unemployment Insurance Laws, Effective January 2014.

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. First, our design allows us to compare outcomes for respondents who receive regressive versus progressive framings of the same policy. This is a key distinction from the vignettes 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 presenting any distributional information.

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. 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.

Figure 2: Experiment 2 Distributive Treatment Results



Notes: Panels show average differences between regressive and progressive treatments for the HMID and UI policy arms. Panels *a* and *b* show levels of support (spending on program) 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. Panels *c* and *d* show other related outcomes by treatment: “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 should get more benefits, higher-income households should get less benefits (= 1), benefits should be kept about the same (= 0), higher-income households should get more benefits, lower-income households should get less benefits (= -1); “Change Premiums” are the avg. responses to whether lower-income individuals should pay a lower premium, higher-income individuals should pay a higher premium (= 1), premiums should be kept about the same (= 0), lower-income individuals should pay higher premium, higher-income individuals should pay lower premium; “Inequality” are average responses to whether the policy increases econ. ineq. (= -1), keeps econ. ineq. about the same (= 0), decreases econ. ineq. (= 1). Panels *a* and *b* show 95% confidence intervals; panels *c* and *d* report p-values from Pearson’s χ^2 and two-sided, two-sample *t*-tests. P-values are 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.

Panels a and b of Figure 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 either policy. It is worth highlighting that this support response, recorded on a 5-point scale, is coded between -1 and 1. 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. To make sure that respondents understood the distributional information presented to them, for each policy we ask how the policy affects economic inequality. In both policy arms, respondents who saw the regressive frame believe that the policy increases economic inequality, on average, compared to respondents who saw progressive frame (see row 3 of Figure 2c and row 4 of Figure 2d). Therefore, despite perceiving a regressive framing of a policy as increasing economic inequality, respondents do not, on average, appear to substantively change their support when compared to respondents who view the progressive framing.

While citizens may not change their overall support for existing policies after learning about their distributional consequences, they may want to modify how much different income-groups benefit. Indeed, we find that in the case of the HMID, respondents assigned to the regressive treatment are less likely to favor keeping the policy as is (see row 1 of Figure 2c), on average, compared to respondents assigned to the progressive treatment. Furthermore, respondents assigned to the regressive treatment are 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 2 of Figure 2c). These results are partially replicated in the UI policy arm. While respondents favor the status quo at approximately equal levels between treatments, the regressive treatment did cause respondents to favor lower-income households both in terms of increasing benefits and reducing premiums as compared to the progressive treatment, although these differences do not reach conventional levels of statistical significance ($p_{benefit} = 0.25$, $p_{premiums} = 0.08$, see rows 3–4 of Figure 2d).

Contrary to previous studies, we therefore show that citizens do not desire to eliminate policies

that are regressive either in absolute or relative terms. Regressive framings of the HMID and UI system did not, on average, cause respondents to support decreasing government spending on these programs compared to the progressive framings. However, respondents do want to make benefits and costs more progressive when the policy is framed regressively. While this instinct to change the status quo is consistent with the literature, it is important to note that this finding is not limited to the HMID, an indirect real-world policy. We find some similar evidence in relation to the UI system as well. This is significant as the UI system is a highly visible direct transfer from the government, and thus in light of prior research we might expect respondents to be less swayed by being given information about the distributional consequences of the program.

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 leverage the second manipulation in our second study, the indirect/direct treatments discussed in the previous section, to see if these framings affect support for the HMID. As before, we ask respondents if they prefer to increase or decrease benefits provided under the policy, with an option to keep levels about the same.

We field this design on two different samples. The first sample consists of the same SSI respondents used in the sections above. In the second sample, we recruited participants on Amazon’s Mechanical Turk (MTurk).¹⁵ 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

¹⁵For a comparison between MTurk, student convenience, and ANES samples, see Berinsky, Huber and Lenz (2012).

Table 3: Direct v. Indirect Framing of HMID

	\widehat{ATE}_{SSI}	p	\widehat{ATE}_{MTurk}	p
Support	0.04	0.15	0.09	0.02
n	1,241		1,066	

Notes: Table shows average differences between direct and indirect treatments (ATE). “Support” (spending on program) 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.

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. The complete MTurk distributional treatments are shown in the Online Appendix. To analyze the MTurk sample responses, we regress the response on the direct/indirect treatment indicator and indicators that correspond to the distributional treatments. The coefficients on the indirect treatment indicator are reported in Table 3.

We find different results in the different samples (see row 1 of Table 3). In the SSI sample, respondents state a small but statistically insignificant preference for the indirect framing of the HMID relative to the direct treatment ($\widehat{ATE} = 0.04$ 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 ($\widehat{ATE} = 0.09$ in increased support, $p = 0.02$). 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.¹⁶

¹⁶The 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

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.¹⁷

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. Importantly, we show that this preference exists despite evidence that citizens believe that indirect delivery would lead to less take-up among eligible individuals relative to equivalent direct transfers. Additionally, we show that in the case of job training, respondents perceive that indirect delivery of benefits would have regressive distributional effects.

Second, we test whether providing citizens information regarding the distributive consequences of existing policies—utilizing both regressive 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,

interact with the direct/indirect treatment in ways that are different than we observe with SSI respondents.

¹⁷We 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.

both in absolute terms and in comparison to the progressive framing. We also replicate this result in an analogous experiment that measures citizen support for the UI system. While we do not observe changes in overall support for existing policies when citizens learn about their distributional consequences, we document a desire to make the benefits and costs of both government programs more progressive when they are framed regressively.

Our findings have several implications that advance the study of policy design and policy feedbacks. We show that citizens prefer indirect government spending notwithstanding their belief that such transfers lead to fewer beneficiaries and greater upward bias (in the case of job training) as compared with direct spending measures. 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 welfare 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. We suggest that future work into the motivations that undergird citizen preferences for indirect spending reorient their investigation towards perceptions of the intended targets of prospective government spending. One such area that we believe is ripe for inquiry are attitudes regarding citizen dependency on government spending.

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 suggests 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.¹⁸

¹⁸A 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=.e57003cdc7.

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 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.

REFERENCES

- Arnold, R. Douglas. 1992. *The Logic of Congressional Action*. Yale University Press.
- Berinsky, Adam J., Gregory A. Huber and Gabriel S. Lenz. 2012. "Evaluating Online Labor Markets for Experimental Research: Amazon.com's Mechanical Turk." *Political Analysis* 20(03):351–368.
- Berinsky, Adam J., Michele F. Margolis and Michael W. Sances. 2014. "Separating the Shirkers from the Workers? Making Sure Respondents Pay Attention on Self-Administered Surveys." *American Journal of Political Science* 58(3):739–753.
- Campbell, Andrea 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 University Press.
- Charness, Gary, Uri Gneezy and Michael A. Kuhn. 2012. "Experimental methods: Between-subject and within-subject design." *Journal of Economic Behavior and Organization* 81(1):1–8.
- Dunning, Thad. 2012. *Natural Experiments in the Social Sciences: A Design-Based Approach*. Strategies for Social Inquiry Cambridge University Press.
- Ellis, Christopher and James A. Stimson. 2012. *Ideology in America*. Ideology in America Cambridge University Press.
- Faricy, Christopher and Christopher Ellis. 2013. "Public Attitudes Toward Social Spending in the United States: The Differences Between Direct Spending and Tax Expenditures." *Political Behavior* 36(1):53–76.
- Gilens, Martin. 2009. *Why Americans Hate Welfare: Race, Media, and the Politics of Antipoverty Policy*. Studies in Communication, Media, and Public Opinion University of Chicago Press.

- Hacker, Jacob S. and Paul Pierson. 2010. *Winner-Take-All Politics*. How Washington Made the Rich Richer—and Turned Its Back on the Middle Class Simon and Schuster.
- Hacker, Jacob S. and Paul Pierson. 2017. *American Amnesia*. How the War on Government Led Us to Forget What Made America Prosper Simon and Schuster.
- Haselswerdt, Jake and Brandon L. Bartels. 2015. “Public Opinion, Policy Tools, and the Status Quo: Evidence from a Survey Experiment.” *Political Research Quarterly* 68(3):607–621.
- Howard, Christopher. 2007. *The Welfare State Nobody Knows: Debunking Myths about U.S. Social Policy*. Princeton University Press.
- Jacoby, William G. 1994. “Public Attitudes toward Government Spending.” *American Journal of Political Science* 38(2):336.
- Kuziemko, Ilyana, Michael I. Norton, Emmanuel Saez and Stefanie Stantcheva. 2015. “How Elastic Are Preferences for Redistribution? Evidence from Randomized Survey Experiments †.” *The American Economic Review* 105(4):1478–1508.
- Malhotra, Neil and Yotam Margalit. 2010. “Short-Term Communication Effects or Longstanding Dispositions? The Public’s Response to the Financial Crisis of 2008.” *The Journal of Politics* 72(3):852–867.
- Malhotra, Neil, Yotam Margalit and Cecilia Hyunjung Mo. 2013. “Economic Explanations for Opposition to Immigration: Distinguishing between Prevalence and Conditional Impact.” *American Journal of Political Science* 57(2):391–410.
- Mettler, Suzanne. 2011. *The submerged state: How invisible government policies undermine American democracy*. University of Chicago Press.
- Morgan, Kimberly J. and Andrea L. Campbell. 2011. *The Delegated Welfare State: Medicare, Markets, and the Governance of Social Policy*. Studies in Postwar American Political Development Oxford University Press.

ONLINE APPENDIX A. SUPPORTING INFORMATION

Table A.1: SSI Sample, Experiment 1:
Respondent Characteristics by Nutrition Treatment Assignment

Covariate	Direct	Indirect	p-value
Age (years)	45.26	44.77	0.40
Black (1 = yes)	0.08	0.08	0.79
Latino (1 = yes)	0.07	0.08	0.35
Democrat (1 = yes)	0.41	0.43	0.36
Republican (1 = yes)	0.38	0.35	0.17
College (1 = yes)	0.47	0.48	0.72
Income (see notes)	69.71	72.78	0.16
n / F test	1,123	1,141	0.69

Notes: Independents who say they are closer to the Democratic (Republican) Party are coded as such. Income is coded as the midpoint of household income intervals selected by the respondent in the pre-survey. Means and proportions are shown; We report p-values from 2-sample t and Pearson χ^2 tests. Treatment assignment is regressed on age, indicators for minority identification, party ID, and college degree, as well as income-group fixed effects to test for the joint orthogonality of pre-treatment covariates and treatment assignment—we report the p-value from the corresponding F test.

Table A.2: SSI Sample, Experiment 1:
Respondent Characteristics by Job Training Treatment Assignment

Covariate	Direct	Indirect	p-value
Age (years)	44.77	45.26	0.40
Black (1 = yes)	0.08	0.08	0.79
Latino (1 = yes)	0.08	0.07	0.35
Democrat (1 = yes)	0.43	0.41	0.36
Republican (1 = yes)	0.35	0.38	0.17
College (1 = yes)	0.48	0.47	0.72
Income (see notes)	72.78	69.71	0.16
n / F test	1,141	1,123	0.69

See Table A.1 notes.

Table A.3: SSI Sample, Experiment 2:
Respondent Characteristics by HMID Distribution Treatment Assignment

Covariate	Regressive	Progressive	p-value
Age (years)	44.63	45.41	0.32
Black (1 = yes)	0.07	0.08	0.80
Latino (1 = yes)	0.07	0.08	0.57
Democrat (1 = yes)	0.43	0.40	0.34
Republican (1 = yes)	0.37	0.38	0.85
College (1 = yes)	0.47	0.48	0.81
Income (see notes)	71.28	72.04	0.80
n / F test	632	609	0.70

See Table A.1 notes.

Table A.4: SSI Sample, Experiment 2:
Respondent Characteristics by UI Treatment Assignment

Covariate	Regressive	Progressive	p-value
Age (years)	46.07	44.66	0.15
Black (1 = yes)	0.07	0.09	0.28
Latino (1 = yes)	0.08	0.07	1.00
Democrat (1 = yes)	0.40	0.46	0.14
Republican (1 = yes)	0.40	0.34	0.08
College (1 = yes)	0.46	0.47	0.82
Income (see notes)	75.62	68.33	0.06
n / F test	415	394	0.35

See Table A.1 notes.

Table A.5: SSI Sample, Experiment 2:
Respondent Characteristics by HMID Directness Treatment Assignment

Covariate	Direct	Indirect	p-value
Age (years)	45.71	44.36	0.09
Black (1 = yes)	0.09	0.06	0.03
Latino (1 = yes)	0.08	0.07	0.37
Democrat (1 = yes)	0.40	0.42	0.38
Republican (1 = yes)	0.39	0.36	0.26
College (1 = yes)	0.47	0.47	1.00
Income (see notes)	71.43	71.90	0.87
n / F test	611	630	0.18

See Table A.1 notes.

Table A.6: MTurk Sample:
Respondent Characteristics by HMID Directness Treatment Assignment

Covariate	Direct	Indirect	p-value
Age (years)	36.56	36.56	0.99
Black (1 = yes)	0.08	0.07	0.70
Latino (1 = yes)	0.04	0.04	1.00
Democrat (1 = yes)	0.52	0.57	0.11
Republican (1 = yes)	0.29	0.28	0.78
College (1 = yes)	0.54	0.54	1.00
Income (see notes)	36.19	34.52	0.36
n / F test	531	535	0.18

See Table A.1 notes.

Figure A.1: Distributional Information Treatments Employed on MTurk Sample

(a) Regressive Framing

The table below shows the average benefit amounts eligible households of different incomes receive from the Home Mortgage Interest Deduction in a year. For example, an eligible household earning less than \$20,000 a year gets an average benefit of \$275.

Annual Household Income	Avg. Benefit Received/Year
Below \$20,000	\$275
\$20,000 - \$39,999	\$552
\$40,000 - \$74,999	\$930
\$75,000 - \$99,999	\$1,244
\$100,000 - \$199,999	\$2,067

(b) Progressive Framing

The table below shows the average benefit amounts eligible households of different incomes receive from the Home Mortgage Interest Deduction in a year as a percentage of the taxes they pay. For example, an eligible household earning less than \$20,000 a year gets an average benefit of \$36 for every \$100 in taxes paid.

Annual Household Income	Avg. Benefit Received as % of Taxes Paid/Year
Below \$20,000	36%
\$20,000 - \$39,999	23%
\$40,000 - \$74,999	17%
\$75,000 - \$99,999	13%
\$100,000 - \$199,999	11%

(c) Regressive and Progressive Framing

The table below shows the average benefit amounts eligible households of different incomes receive from the Home Mortgage Interest Deduction in a year. For example, an eligible household earning less than \$20,000 a year gets an average benefit of \$275, getting back an average of \$36 for every \$100 in taxes paid.

Annual Household Income	Avg. Benefit Received/Year	Avg. Benefit Received as % of Taxes Paid/Year
Below \$20,000	\$275	36%
\$20,000 - \$39,999	\$552	23%
\$40,000 - \$74,999	\$930	17%
\$75,000 - \$99,999	\$1,244	13%
\$100,000 - \$199,999	\$2,067	11%