Narrative Premiums in Policy Persuasion

Keisuke Kawata[*] Kenneth McElwain[†] Masaki Nakabayashi[‡]

November 23, 2022

Abstract

Survey experiments have shown mixed results about the effect of information provision on attitudes towards controversial policies. We argue that one reason is varied receptiveness to different modes of information. Prior research suggests that people selectively ignore factual, statistical information that contradicts prior beliefs, but are more attentive to narrative information that describes individual experiences. We test this in the context of Japanese attitudes towards poverty relief programs, which tend to be less popular than other welfare expenditures. Using a conjoint survey, we show that there is a "narrative premium": respondents who are shown a narrative story about the plight of a single mother are more likely to support higher expenditures on poverty relief than those who are shown statistical information about the share of single parents living in poverty. This premium is particularly effective in strengthening the convictions of those who are already aware of levels of societal poverty.

Keywords Conjoint experiment; persuasion; communication; poverty relief; narrative-evidence treatment; statistical-evidence treatment; multidimensional policy space; policy vector; Japan.

JEL classification codes D19; H53; I38.

[*]Institute of Social Science, The University of Tokyo. E-mail: keisukekawata@iss.u-tokyo.ac.jp
[†]Institute of Social Science, The University of Tokyo. E-mail: mcelwain@iss.u-tokyo.ac.jp
[‡]Institute of Social Science, The University of Tokyo. E-mail: mn@iss.u-tokyo.ac.jp

1 Introduction

Democratic governments have greatly expanded their provision of public goods since the 20th century. Public education systems have improved labor force productivity, public works projects have reduced transportation and communication costs, and public health insurance has ameliorated the individual risks and costs of illnesses. However, one program that remains controversial is poverty relief, such as direct cash transfers and food assistance to low-income families.[1] While few politicians or voters openly support destitution, programs designed to alleviate poverty are disputed on practical and normative grounds. First, poverty relief may be seen as distorting market incentives by disincentivizing labor (Krueger and Meyer, 2002; Tatsiramos and van Ours, 2014). Second, there may be concerns about people trying to cheat the system to receive unwarranted benefits, or for politicians to manipulate funds to reward supporters (Diaz-Cayeros et al., 2016). Third, because poverty relief is inherently targeted to poorer citizens, perceptions of who "deserve" help become intertwined with prior beliefs about those with lower socioeconomic status. This includes stereotypes of "black welfare queens" in the United States (Gilens, 1999) or of immigrants who take advantage of generous welfare provisions in wealthier countries (Barrett and McCarthy, 2008; Fix and Haskins, 2002; Øland, 2019; Valentino et al., 2002). These stereotypes, in turn, influence the very design of poverty relief programs, including stringent means-testing and political shaming of the poor.

One issue that may drive such apprehensions is misperceptions of the demographic profiles of beneficiaries. While the causes of poverty are complex, one group that comprises a large share of the poor is single-parent households with small children (Maldonado and Nieuwenhuis, 2015). This is a subpopulation for which many citizens have some sympathy,

[1]Depending on national and temporal context, poverty relief is alternatively called poverty reduction or poverty alleviation. For simplicity, we refer to these programs as "poverty relief" throughout the paper.

and awareness of their plight may increase public support for expanding poverty relief. However, correcting public misperceptions is easier said than done. While the provision of "accurate" information may update beliefs, there is mixed evidence about the effect of factual corrections on people's preferences (Hopkins et al., 2019). Nyhan (2020) attributes this to people's doubts about the trustworthiness of the information and their tendency to reject evidence or arguments that run counter to prior beliefs. This type of motivated reasoning is particularly strong for politically controversial issues where partisan claims are built on emotional, as well as informational, appeals.

This is not to say that beliefs are inelastic. A vibrant literature in the social sciences and public health has examined whether the persuasiveness of information changes with its *mode of presentation*. One notable distinction is between "statistical" information, which distills individual behavior into one aggregate metric, and "narrative" information, which uses individual cases to illuminate broader social patterns. In a survey of the state of the field, Winterbottom et al. (2008) argues that narrative evidence influences decision-making more than statistical evidence. Many studies similarly note the greater power of narrative information, although there are disagreements about why and to what extent. In public health, De Wit et al. (2008) demonstrates that narrative evidence is more effective than statistical evidence in making recipients aware of infection risks. In economics, Akerlof and Snower (2016), Chater and Loewenstein (2016), and Shiller (2017) emphasize the role of narratives in decision making, although they do not identify which elements of narratives are more persuasive than statistical evidence. Other studies find that the difference between the two modes are slight (Mazor et al., 2007; Hong and Park, 2012; Guo et al., 2019), while yet others argue that the key is to combine both narrative and statistical evidence (Nan et al., 2015; Ferraro et al., 2011; Ferraro and Price, 2013).

This paper draws on these studies to examine how the presentational mode of information can shape individual preferences about appropriate levels of poverty relief funding. In doing so, we aim to contribute to the literatures on welfare politics, social psychology, and comunication. Our analysis is based on a conjoint survey experiment conducted in Japan. The pursuit of a broad middle-class society (*sōchūryū shakai*) has long been considered a core national value in Japan, but poverty rates have increased rapidly since the 1990s, making it on par with most developed democracies (Shirahase, 2014; Chiavacci, 2008). While debates about poverty and government programs cannot be fully separated from national context, Japanese discourse has not hinged on parallel debates about race or immigration, making its results easier to generalize across countries. We discuss the Japanese context in greater detail in later sections.

Our experiment randomly assigns respondents to one of three groups, which include one control and two treatment arms. Each treatment group is shown information about single-parent households living in poverty. The "narrative" treatment is a short essay about a single mother who works full-time in a low-wage job and depends on food banks to feed her child. The "statistical" treatment shows factual information that a majority of single-parent households live in relative poverty.

Following these treatments, respondents are asked to compare different budget proposals in a conjoint design that randomly varies the expenditure levels of five fiscal programs, including poverty relief, with debt repayment as the alternative. The conjoint design is a critical element of this paper's contributions. Past studies of poverty relief have largely looked at preferences on a unidimensional space, such as more or less spending on a specific program, but this design makes it difficult to separate attitudes towards poverty relief from prior beliefs about the merits of big versus small government. Our conjoint design explicitly incorporates the multidimensional nature of budget proposals. This allow us to examine the tradeoffs that respondents make in spending across policy vectors, as well as their preference for an aggregate increase in expenditures as opposed to fiscal retrenchment via debt repayment.

Our quantity of interest is the "narrative premium", or how much more effective narrative evidence is than statistical evidence in changing preferences about fiscal allocations across multidimensional policy space. We estimate differences in the marginal means, or the average change in the rating of a budget proposal caused by each spending level of each public program. We find that the narrative premium increases for higher levels of poverty spending: respondents who are shown the narrative about the plight of a single mother are more likely to support an expansion of poverty relief than those who are shown statistical evidence. The narrative premium is statistically significant and positive over all background characteristics we surveyed, including partisanship. Further classification analyses suggest that narrative evidence is effective in strengthening the convictions of those who are already aware of actual levels of societal poverty, rather than converting those who underestimate the issue.

The remainder of the paper is organized as follows. Section 2 introduces the narrative premium and provides case context about Japan, including recent trends in the poverty rate and fiscal constraints. Section 3 explains our information treatment, survey structure, and the randomized conjoint experiment. Section 4 discusses our identification strategy. Section 5 reports the results of our conjoint analysis and the effect of our information treatments. Section 6 concludes the paper.

2 Narrative Premium

2.1 Perceptions of Information

A variety of social science disciplines have explored how the provision of information can alter citizens' perceptions about the merits of public policies and state interventions. A major point of contention is the relative value of different *types* of information, broadly categorized into statistical versus narrative evidence. On the one hand, statistical data about aggregate patterns in society distill the preferences or behavior of a large number of people into one metric. It may be seen as more objective and credible, which results in greater message acceptance by the receiver. On the other hand, narrative information, such as case stories about the conditions of specific individuals, may elicit a stronger emotional response and generate greater identification with the message. Evidence about the relative persuasiveness of these information modes is mixed. In their meta-analysis of experimental studies that use statistical versus narrative evidence, Allen and Preiss (1997) find that statistical evidence is more persuasive. By contrast, De Wit et al. (2008) examine perceptions of personal health risk from hepatitis B and show that anecdotal personal accounts were more likely to raise perceptions of risk and intentions to obtain vaccinations than statistical evidence about the prevalence of hepatitis B.

One reason why narrative evidence may be more powerful is that they are less likely to trigger psychological defense mechanisms. Slater and Rouner (1996) argue that statistical evidence may be more persuasive when the message is consistent with prior preferences, but that narrative evidence is more persuasive when the message requires recipients to change strongly held beliefs. Kalla and Broockman (2020) examine how exclusionary attitudes, or prejudice against outgroups and opposition to policies that help them, may be mitigated. Drawing on the psychology literature, they contend that people do not want to admit their views are in error, which threatens their sense of autonomy. However, interpersonal, non-judgmental conversations or exchanges of narratives can reduce exclusionary attitudes durably.

Another mechanism is related to the importance of stimulating emotional responses. The availability heuristic, proposed by Tversky and Kahneman (1973), suggests that people have an easier time recalling information that is presented vividly. Relatedly, narrative information may be more likely to evoke strong emotional reactions that trigger behavior independent of cognitive judgements (Loewenstein et al., 2001). Brader et al. (2008) look at this emotional dimension explicitly in the context of immigration policy. While evidence about the costs of immigration increase opposition to its expansion, this is largely mediated by emotions of anxiety, which are more likely to be triggered among white respondents when stories about Latino, not European, immigrants are featured.

2.2 Fiscal Tradeoffs and Case Context

The relative persuasiveness of different information modes also applies to social attitudes towards public policies, particularly those with low baseline popularity levels. Government budgets are the culmination of myriad choices about how much to spend where and on whom. The redistributive nature of fiscal decisions is also relevant to programmatic spending, including those that are often described as public goods, such as healthcare, unemployment, and poverty relief (Estevez-Abe, 2008). For example, health care systems may be less valued by the young and healthy, and public education may be less valued by those who attend (or send their children to) private schools. Most social insurance programs are about mitigating risk, but the risks of illness may be more random than the risks of being injured on the job, generating competing preferences about mandating public coverage (Mares, 2003). In other words, not all programs are viewed positively, either because one does not expect to qualify for those benefits, or because the beneficiaries are somehow undeserving of public support.

One such example is poverty relief programs. Actual levels of poverty have increased in many advanced industrialized economies since the 1980s, but funding for poverty relief programs has failed to keep pace, due in part to neoliberal turn in the 1990s that emphasized market-based solutions to social problems. The poor also tend to vote at lower rates and are less likely to be mobilized by political entrepreneurs, albeit with cross-national and temporal differences based on electoral geography (Jusko, 2017). An illuminating case to explore is Japan. Income distribution was drastically de-concentrated during WWII (Moriguchi and Saez, 2008), and Japanese citizens have long considered their society to be an equal one. However, inequality has gradually risen since the 1990s, eroding people's impressions of Japan as an egalitarian society (Chiavacci, 2008; Hommerich and Kikkawa, 2019; Kanbayashi, 2019).

Figure 1 shows 2015 data on poverty rates among the seven major advanced economies (G7). Japan's relative poverty rate of 15.7% (triangle marker) is the second highest after the US. Multiple causes have contributed to growing income inequality and poverty. One is the transformation of wage determination toward more competitive performance-based compensation. Another is the rise in the share of workers under part-time and limited-duration contracts, whose wages have barely increased since the 1990s. These factors collectively squeezed the income of the middle class. Rapid aging has also raised poverty rates in Japan (Shirahase, 2021; Tachibanaki, 2006), with those over 65 years of age currently registering at 19.6% (diamond). Of particular concern is the high poverty among those aged 0-17, which is currently at 13.9% (circle). Parental poverty due to the structural changes described above has already adversely affected child health (Kohara et al., 2019; Nakamura, 2014).

Poverty rate, percent of total population, 2015

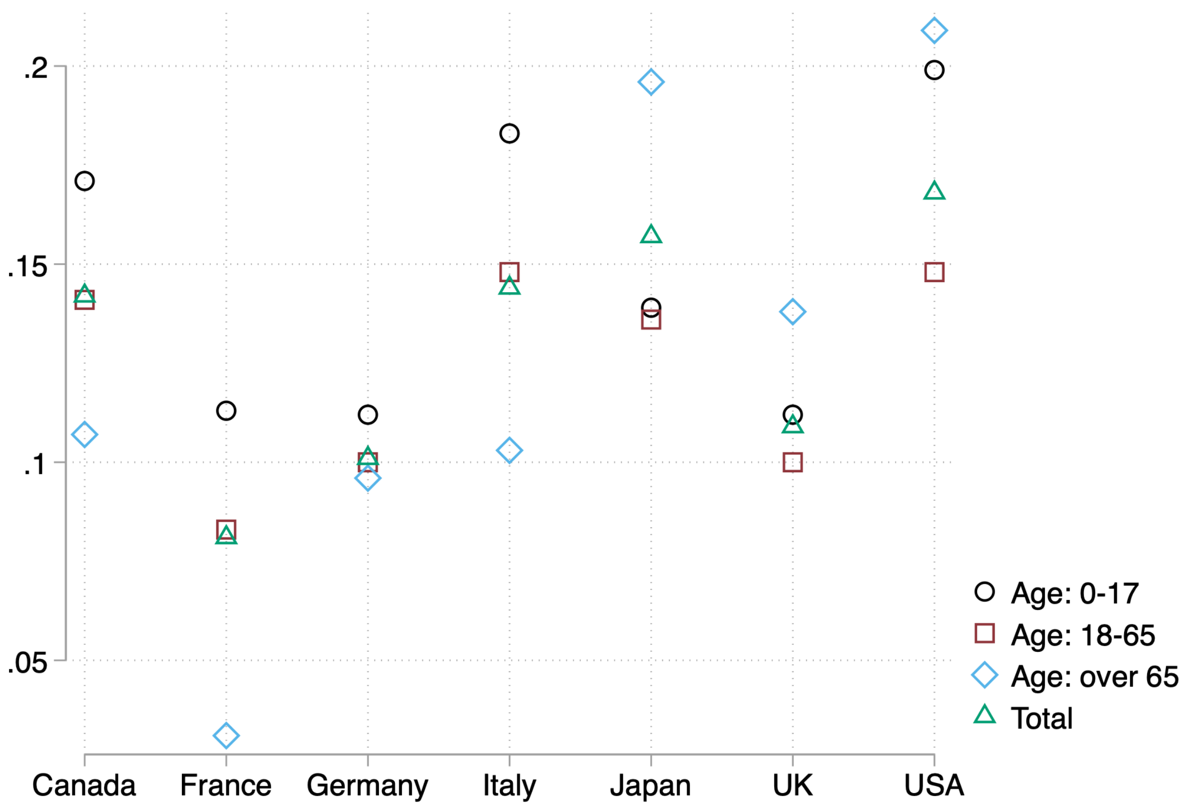

Figure 1: Source: OECD (https://doi.org/10.1787/0fe1315d-en; accessed on May 23, 2021.

While the Japanese state has a robust welfare system that includes national health insurance and a national pension system for all residents, its public assistance component is less generous. Article 25 of the Constitution of Japan guarantees that "All people shall have the right to maintain the minimum standards of wholesome and cultured living" and requires the State to promote social welfare, security, and public health. The Public Assistance Act of 1950 was passed in response to this obligation, but the government has interpreted what constitutes "minimum standards" flexibly. Notably, to be eligible for public assistance, individuals must first demonstrate that they cannot receive sufficient support from stem family members.[2] In practice, a family is considered for eligibility if its income is below the relative poverty line (Inaba, 2011). Once eligible, the State pays living expenses, assistance benefits for rent, education expenses, medical expenses, and long-term care expenses, although the amounts fall short of satisfactory poverty relief.[3] As of 2019, there were 1.64 million recipient families and 2.07 million individual recipients.[4] This assistance comes with strings. The law requires recipients to work or search for a job during the period of benefit. Also, recipients are obliged to follow instructions from the municipal welfare office in charge of the family.

Solutions to poverty include increasing the generosity of public assistance and expanding eligibility, but these run into conflict with competing fiscal demands. For one, demographic aging has steadily increased health care and pension costs, even as labor force size has declined. The share of social security contributions in GDP has risen from approximately 7.5% in 1990 to 12.5% in 2017, making it on par with welfare states in continental Europe. [5] Its share in total tax contributions is the highest within the G7. This expansion of Japan's welfare state has been financed by ever-expanding government debt, which began to rise in the mid-1990s and has been well beyond 200 percent of GDP in the last decade.

How to address rising poverty in the face of demographic aging and rising welfare costs is a fundamental dilemma in most advanced-industrialized democracies. While poverty is

[2] The Civil Code of 1896 stipulates the duty of stem family members. The Public Assistance Act of 1950 is exercised (Article 4) only if support from family members falls short of the "minimum standard of wholesome and cultured living" stipulated by Article 25. http://www.japaneselawtranslation.go.jp/law/detail_main?re=&vm=02&id=24; accessed on March 18, 2021.

[3] As of 2020, if the family is comprised of a husband, wife, and one children and lives in Tokyo, living expenses assistance only amounts to JPY158,760 per month. If the family is composed of a mother and two children, it amounts to JPY190,550. These amounts draw on model cases provided by the Ministry of Health, Labour and Welfare: https://www.mhlw.go.jp/content/000578652.pdf; accessed on March 18. 2021. The latter case implies 2.29 million yen per year, which would place this model family in the bottom 25th percentile, per an MHLW household income survey in 2019 (see Table A1 in the online appendix).

[4] https://www.mhlw.go.jp/topics/2020/01/dl/9_shakaiengo-03.pdf; accessed on March 2021.

[5] (OECD) https://data.oecd.org/; accessed on July 20, 2020.

no longer an ignorable issue in Japan, augmenting relief programs may necessitate further deficit spending, tax increases, and/or cuts in other fiscal expenditures, all of which may provoke political backlash. However, governments may be more likely to support poverty relief if its expansion—and any associated tradeoffs—are accepted by voters. While poverty relief is a small fraction of all government spending, its perceived necessity and efficacy is shaped by prior political beliefs and social stereotypes. These can moderate the willingness of citizens to be persuaded of the need to expand programmatic spending.

2.3 Assessing Attitudes Toward Poverty Relief

Political scientists have explored the salience of motivated reasoning—ignoring information that contradicts one's beliefs—in the context of public policies aimed at helping the poor. Many point to the ability of narrative information to change prior preferences. For example, there is substantial evidence that policy choices may be framed in ways that activate negative stereotypes of program beneficiaries. Gilens (1996) argues that means-tested transfer programs, often termed "welfare", can stimulate negative views of blacks among white Americans. Manipulating information about the race of welfare recipients experimentally, he finds that prior biases, specifically negative views of black welfare recipients as lazy, generates greater opposition to welfare than do negative views of white recipients. However, narrative evidence that discounts this bias can persuade respondents to support welfare programs. At the same time, people's racial priors are not inherently negative. Chudy (2021) shows that whites with racial sympathy are more likely to support government aid targeted to blacks. We should note here that the generalizability of American patterns to other nations should be approached carefully, since not all countries have ingrained tensions between ethnic or religious minority groups. However, existing work points consistently to a tendency for poverty relief programs to be seen as benefiting "unworthy" people.

That said, we believe that there are two limitations to existing studies, one relating to treatment design, and one relating to the estimated quantity of interest. First, most experimental studies on fiscal choices typically look only at the impact of narrative information. This may be because existing debates are intimately intertwined with race and politics, where psychological biases or stereotypes are ingrained deeply. To better understand the persuasiveness of information, however, it is necessary to compare the effects of narrative versus statistical information directly. Second, most studies look at changes in attitudes to specific policy items on one dimension, such as whether to spend more or less on poverty relief. However, there are opportunity costs to expanding fiscal programs. The real choice facing citizens is not simply whether poverty relief should be increased, but whether that is evaluated more highly than changing expenditures on other programs, deficit spending levels, and tax rates.

With this in mind, this paper adopts an experimental design wherein we randomly manipulate narrative versus statistical information to estimate their relative effects explicitly. We then implement a conjoint analysis, where respondents are asked to evaluate randomized budget allocations across five policy areas, including poverty relief, relative to paying down government debt. In the next section, we explain this experimental design in greater detail.

3 Experimental design

In March 2020, we conducted a non-probability online survey of 15,000 respondents in Japan, recruited by Rakuten Insight. As we will discuss below, the demographic characteristics of our respondents do not vary significantly from the Japanese population. Further details on our research ethics protocol and sampling strategy can be found in Online Appendix B.

3.1 Conjoint analysis

Our outcome of interest is respondents' preferences on how to allocate additional revenue from an increase in the consumption tax (analogous to a value-added tax) rate from 8 percent to 10 percent. This tax increase was implemented in October 2019, after much political debate about its necessity for fiscal balance versus its consequences on consumption levels and growth (Tanaka, 2022).

Respondents were tasked with comparing two hypothetical budgets, which randomly assigned 0, 5, 10, 15, or 20 percent of the additional revenue from the consumption tax hike to (a) "welfare (minimum wages, unemployment benefit, public housing for low-income earners, etc.)," (b) "pensions," (c) "health insurance," (d) investment in "infrastructure (roads, running water, airports, etc.)," and (e) "education (subsidy for tuition, increase in nursery schools, etc.)." They were informed that any residue would be allocated to the redemption of government bonds. Under the Japanese social security system, the National Pension Plan and National Health Insurance are universal insurance policies that cover all residents in Japan. Beneficiaries of increased subsidies to "pensions" and "health insurance" are not limited to poor individuals. Thus, only expenses for (a) "welfare" directly aims an income transfer from rich to poor individuals.

This design is summarized in Table 1. For instance, in one task, one package may assign 10 percent to (a), 5 percent to (b), 5 percent to (c), 20 percent to (d), 20 percent to (e), and 40 percent to the redemption of government bonds (f). The second package may assign 5 percent to (a), 10 percent to (b), 15 percent to (c), 0 percent to (d), 10 percent to (e), and 60 percent to the redemption of government bonds (f). Thus, each package is a menu of how much of the increased tax revenue is to be spent for what, with any savings to be used to reduce government debt. Respondents were then asked to choose the budget package they preferred. This randomized conjoint task was repeated five times, generating a total of 150,000 profile preferences (2 profiles per task $\times$ 5 rounds $\times$ 15000 respondents).

Table 1: Attributes and attribute levels for hypothetical public policies.

Policy attributes	level	Policy attributes	level	Policy attributes	level
Welfare programs	0%	Pensions	0%	Health insurance	0%
	5%		5%		5%
	10%		10%		10%
	15%		15%		15%
	20%		20%		20%
Education	0%	Infrastructure	0%	Residual is to	
	5%		5%	redeem debts	
	10%		10%		
	15%		15%		
	20%		20%		

Let $\boldsymbol{a}_j^{i,r}$ and $\boldsymbol{a}_{-j}^{i,r}$ denote alternative policy packages viewed by respondent i in round r. Also, let $a_{l,j}^{i,r}$ and $a_{l,-j}^{i,r}$ denote the lth attributes of $\boldsymbol{a}_j^{i,r}$ and $\boldsymbol{a}_{-j}^{i,r}$, respectively. Then, $\boldsymbol{a}_j^{i,r}$ and $\boldsymbol{a}_{-j}^{i,r}$ consist of five elements from (a) to (f) described above, such that $a_{l,j}^{i,r}$ and $a_{l,-j}^{i,r}$ take 0, 5, 10, 15, or 20% for $l = 1, ..., 5$, $i = 1, \ldots, 15,000$, and $r = 1, \ldots, 5$. We call a response of i in round r on policy package j a sample. Later in this paper, we classify these samples to groups using sorted group average marginal treatment effects (GATES), as described in Section 4.

3.2 Information treatment

Prior to the conjoint section, respondents were randomly assigned to three groups: narrative information treatment, statistical information treatment, and a control. Both treatments were designed to highlight the financial plight of single-parent households.

Information treatment 1 Those in the narrative treatment group were shown a short essay that described a poor, single-parent family. Given its context, most respondents would understand that this episode is about a single-parent household during the school summer break, when schools do not provide lunch, and that the single mother works as a full-time but non-regular worker, from the description about hourly wages.

> "Noriko works full-time. Thus, Atsushi spends his day at a children's hall. Noriko must make Atsushi take a lunch to the children's hall, different from an elementary school. Treats for snack time are also necessary. Noriko works full-time, but on an hourly wage. Her after-tax income is about 120,000 yen (*Note:* 1,137 US dollars) per month. To make his lunch everyday, she appreciates the support of foodstuff from a food bank. It literally sustains Atsushi's weight." Excerpted from Makoto Yuasa, *"Nantoka Suru" Kodomo no Hinkon ("Do Something" with Child Poverty)*, Tokyo: Kadokawa Shoten, 2017.[6]

Information treatment 2 Those in the statistical treatment group were shown statistical evidence of the relative poverty of single-parent households in Japan as of 2015, as calculated by the Ministry of Health, Labour and Welfare.[7] The pie chart was accompanied by a caption, which wrote "51% of single-parent households, where the parent is working, live in relative poverty." Figure 3.2 is a translated version of the chart.

Out of 15,000 respondents, respondents were randomly assigned to the two treatments and one control, which included no information about poverty. Thus, our sample consists of

[6] Yuasa (2017, pp. 147–148), translated by the authors. The original Japanese text was excerpted in the experiment.

[7] https://www.mhlw.go.jp/toukei/list/dl/20-21-h28_rev2.pdf; accessed on November 30, 2020.

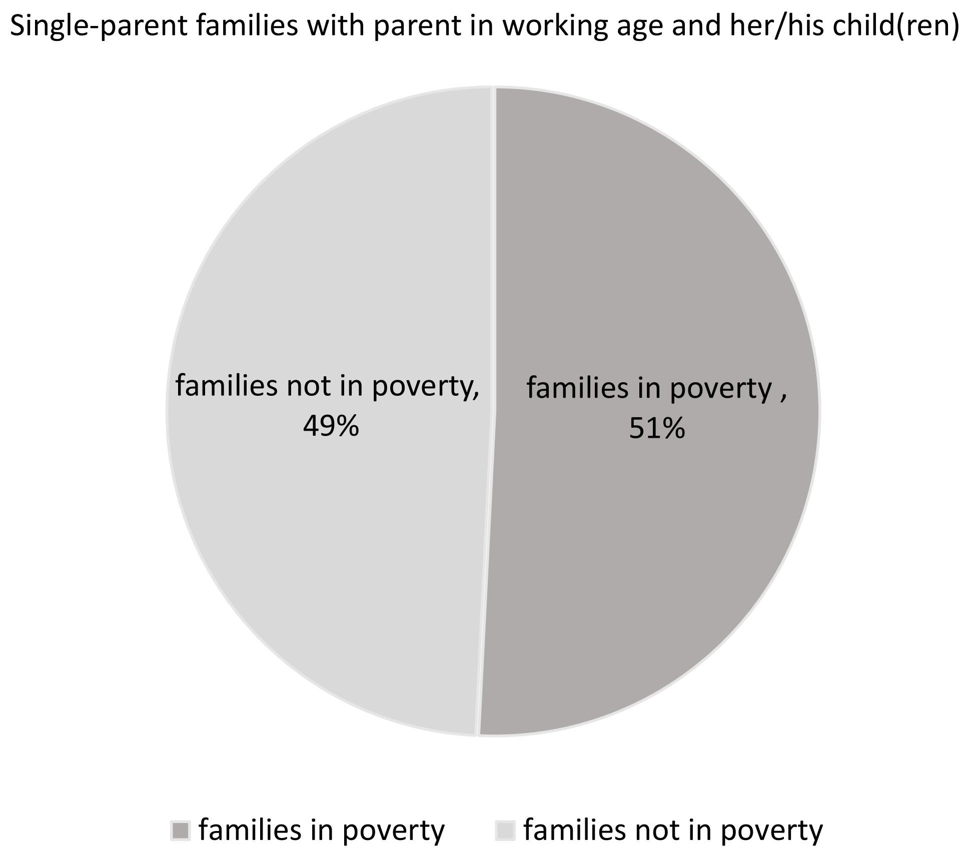

Figure 2: Statistical evidence under information treatment 2.

one-third sub-sample of the narrative evidence-treated, one-third of the statistical evidence-treated, and one-third of the control group.

3.3 Background characteristics and beliefs

In addition to estimating the treatment effect of the two forms of information in our overall sample, we also explore factors that may contribute to treatment effect heterogeneity. We do so using survey items on the background characteristics of respondents. The questions included basic demographic information, such as occupation and education, as well as attitudes relating to politics and society. The information treatment effects may depend on these background traits as well as the budget packages themselves. Thus, as discussed in Section 4, we estimate the information treatment effect as a function of background characteristics as well as the proposed budgets. While explaining these heterogeneities is not a core part of our analysis, we explore patterns using sorted group average marginal treatment effects (GATES), to be discussed below. The primary surveyed characteristics are summarized in Table 2.

Table 2: Primary background characteristics

Category	Characteristics
Demography	· gender/age/prefecture of residence/number of siblings
	· educational background
	· marital status/number of children
	· co-habitation with parents or parents-in-law
Occupation	· working status
	> employed or self-employed
	> regular or non-regular/job title/size of employer
Income	· own income/own household's income
Political preferences	· affinity for a specific party
Perception of poverty	· perceived poverty rate
	· perceived single-parent household poverty rate

Questions on demography include gender; age; prefecture of residence; marital status; number of children; number of siblings; whether the respondent lives with parents, parents-in-law, or neither.[8] On education, we asked for respondents' highest degree of education. Questions on occupation include whether the respondent is at work; if at work, then whether

[8]Prefectures are the main sub-national administrative unit in Japan, of which there are 47.

employed or self-employed; if employed, then whether full-time or part-time, whether as a regular or non-regular worker, job title, and size of the employer. We also asked for respondents' own annual income and their household's annual income.

On political attitudes, we asked for party identification. Esarey et al. (2012), Clark and D'Ambrosio (2015), and Kerschbamer and Müller (2020) demonstrate relationships between the political preferences of respondents and their baseline support for income redistribution policy. Furthermore, Kuziemko et al. (2015) and Alesina et al. (2018) find that political preferences affect not only baseline support for income redistribution, but also information treatment effects on their support for income redistribution policy.

Our survey items on political attitudes addresses such possibilities. Party affinity is an obvious cue: the ruling Liberal Democratic Party (LDP) is center-right in ideological orientation, and its supporters are less likely to favor redistributive policies. In our sample, 26.3% of respondents identified with the LDP, while only 7.6% identified with the center-left Constitutional Democratic Party, the largest opposition party. A 51.8% majority stated that they were independents who did not support any party. An important caveat is that this question was included after the information treatments and conjoint tasks, in order to avoid priming respondents with political cues. One obvious downside is that we cannot fully dismiss the possibility that party identification was influenced by the experimental manipulation. However, we believed this to be a worthwhile trade off for two reasons. First, our main interest was in estimating the treatment effects cleanly, with heterogeneity in such effects being a secondary concern. Second, we do not see any significant difference in party affinity between narrative-evidence, statistical-evidence, and control treated groups, as shown by Figure 3.

Finally, we included two questions regarding perceptions of poverty in Japan: the perceived percentage of households living in poverty generally, and that of single-parent house-

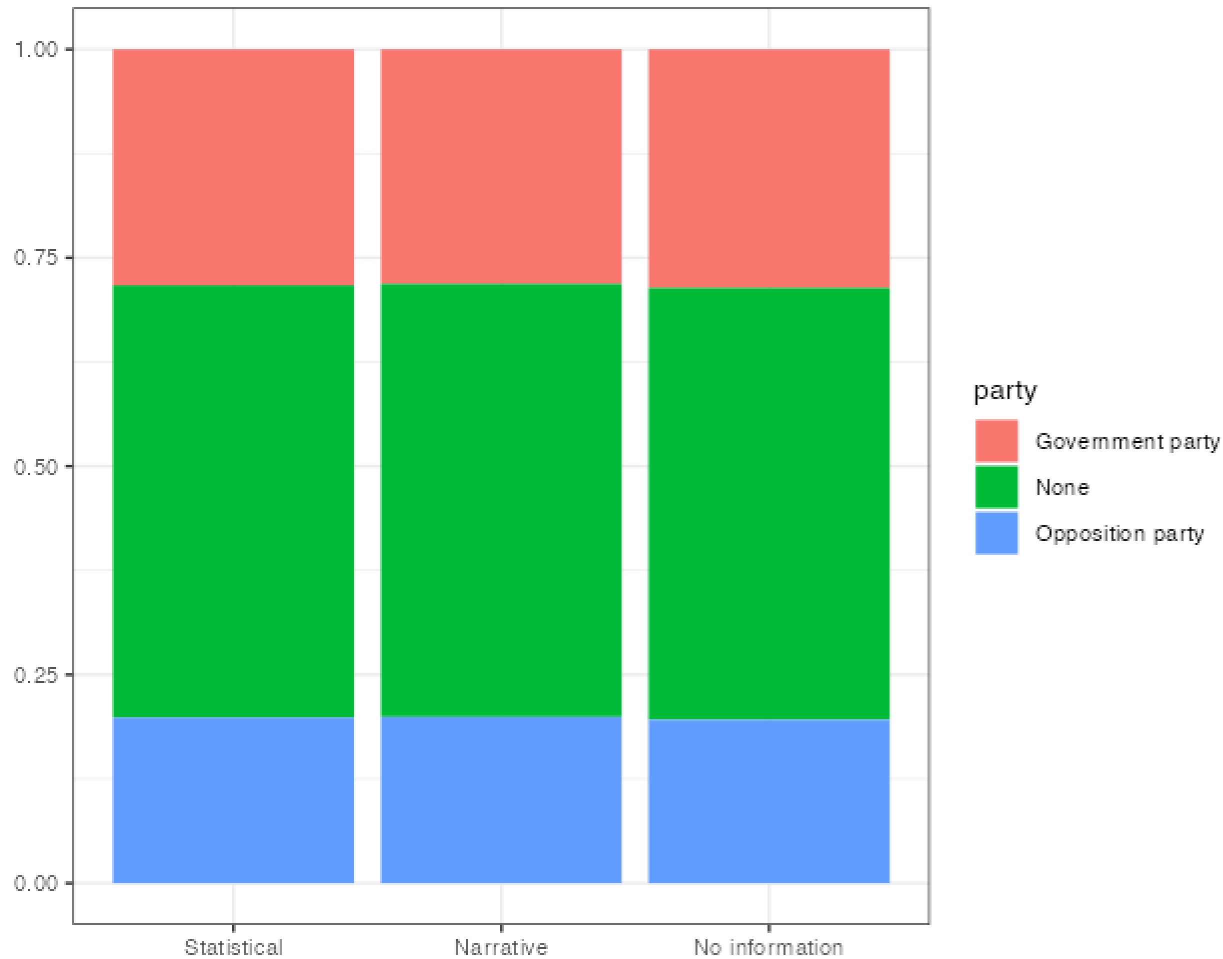

Figure 3: Party approval rates under each treatment arm.

holds specifically. The mean response was 29.9% for the former and 44.2% for the latter. Given actual poverty rates of 15.6% and 50.8%, respectively,[9] it appears that the majority of respondents underestimate the poverty levels of single-parent households. However, a quick caveat is that the distribution of perceived poverty rates was not skewed toward underestimation. While the mean of perceived poverty rates was lower than real levels, the distribution is nearly symmetric in both directions, as shown in Figure 4. Most respondents either overestimated or underestimated child poverty rate of single-parent households and

[9] https://www.mhlw.go.jp/toukei/list/dl/20-21-h28_rev2.pdf; accessed on August 22, 2021.

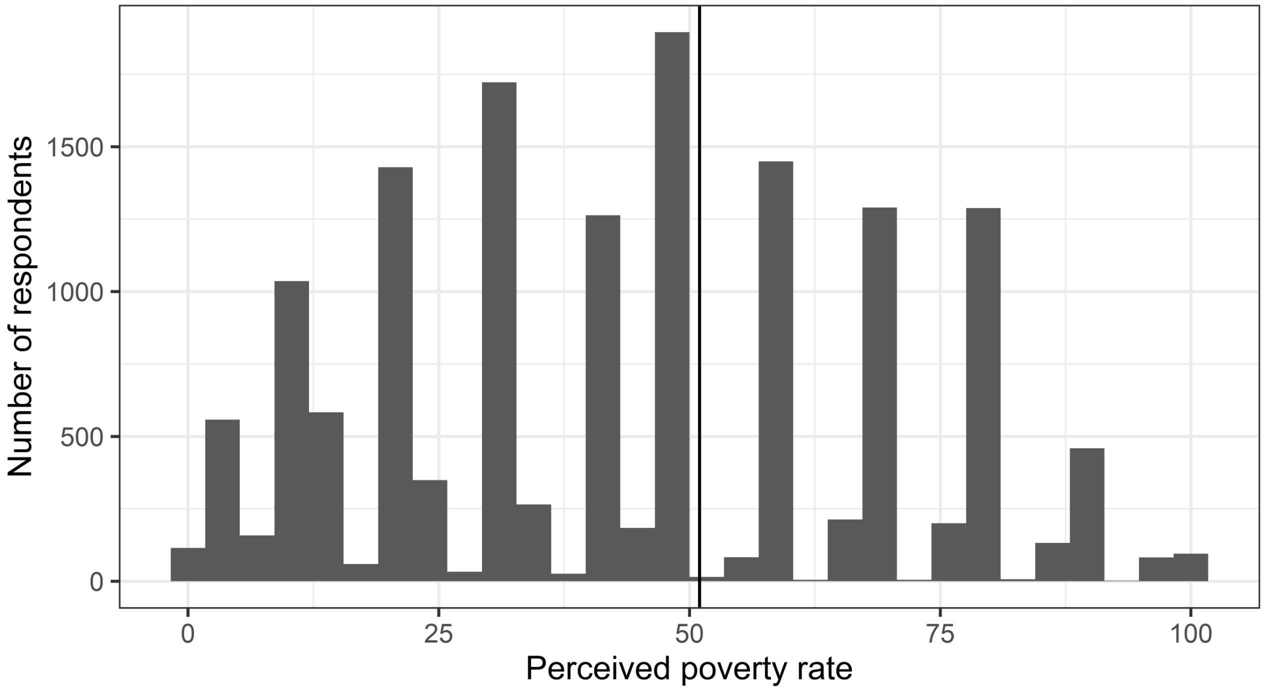

Figure 4: Perceived poverty rate of single-parent households.

did not show a skewed distribution.

Descriptive statistics of the background characteristics are presented in Table 3. Gender takes 1 if the respondent was female and 0 otherwise. Work and marital statuses also take 1 if the respondent worked and were married, respectively. The maximum value of the number of children is 5, such that an answer "5" might include more than five children. Education and income strata are ordinal variables in our statistical models, but in Table 3, each level takes 1 if the respondent selected a particular level and 0 otherwise.

4 Identification Strategy

Consider a public policy space whose dimensions are alternative public policies including fiscal retrenchment. Given the tax revenue, any policy choice is defined as a direction of

Table 3: Descriptive statistics of background characteristics.

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Age	15,000	50.105	14.348	18	39	62	79
Gender	15,000	0.462	0.499	0	0	1	1
Marital status	15,000	0.654	0.476	0	0	1	1
Number of children	15,000	1.171	1.123	0	0	2	5
Education: Junior high school	15,000	0.014	0.119	0	0	0	1
Education: High school	15,000	0.246	0.431	0	0	0	1
Education: Vocational college	15,000	0.115	0.319	0	0	0	1
Education: 2-year college	15,000	0.093	0.291	0	0	0	1
Education: Technical college	15,000	0.014	0.116	0	0	0	1
Education: 4-year college	15,000	0.457	0.498	0	0	1	1
Education: Graduate school	15,000	0.060	0.238	0	0	0	1
Perceived poverty rate %	15,000	29.949	18.266	0	20	40	100
Perceived single-parent household poverty %	15,000	44.211	24.996	0	20	65	100
Working status	15,000	0.693	0.461	0	0	1	1
Own income: Less than JPY 0.5 million	15,000	0.160	0.367	0	0	0	1
Own income: JPY0.5–0.99 million	15,000	0.085	0.279	0	0	0	1
Own income: JPY1–1.49 million	15,000	0.080	0.271	0	0	0	1
Own income: JPY1.5–1.99 million	15,000	0.063	0.242	0	0	0	1
Own income: JPY2–2.49 million	15,000	0.077	0.267	0	0	0	1
Own income: JPY2.5–2.99 million	15,000	0.065	0.246	0	0	0	1
Own income: JPY3–3.99 million	15,000	0.116	0.320	0	0	0	1
Own income: JPY4–4.99 million	15,000	0.100	0.299	0	0	0	1
Own income: Higher than JPY5 million	15,000	0.255	0.436	0	0	1	1
Household income: Less than JPY0.5 million	15,000	0.029	0.168	0	0	0	1
Household income: JPY0.5–0.99 million	15,000	0.013	0.115	0	0	0	1
Household income: JPY1–1.49 million	15,000	0.024	0.153	0	0	0	1
Household income: JPY1.5–1.99 million	15,000	0.036	0.186	0	0	0	1
Household income: JPY2–2.49 million	15,000	0.051	0.220	0	0	0	1
Household income: JPY2.5–2.99 million	15,000	0.053	0.224	0	0	0	1
Household income: JPY3–3.99 million	15,000	0.120	0.324	0	0	0	1
Household income: JPY4–4.99 million	15,000	0.122	0.328	0	0	0	1
Household income: JPY5–5.99 million	15,000	0.116	0.320	0	0	0	1
Household income: JPY6–6.99 million	15,000	0.088	0.283	0	0	0	1
Household income: JPY7–7.99 million	15,000	0.090	0.286	0	0	0	1
Household income: JPY8–8.99 million	15,000	0.065	0.247	0	0	0	1
Household income: JPY9–9.99 million	15,000	0.052	0.222	0	0	0	1
Household income: Higher than 10 million	15,000	0.141	0.348	0	0	0	1

the policy vector in the space.

Let us suppose that treatment effects are heterogeneous over background characteristics. Then we can estimate the treatment effects by the potential outcomes framework (Imbens and Rubin, 2015; Wager and Athey, 2018; Athey and Imbens, 2019). We compare one potential outcome under an information treatment that provides narrative evidence, one under an information treatment that provides statistical evidence, and one without an information treatment. We then identify the difference in the direction of preferred policy vectors between the three potential outcomes as the causal effect of each information treatment.

In our experimental design, one-third each of respondents receives the narrative evidence treatment, the statistical evidence treatment, and no evidence as part of the control. Let $w_i \in \{c, n, s\}$ denote the information treatment indicator, where c refers to the control group, n to the narrative evidence-treated group, and s to the statistical evidence-treated group. Once having received the narrative, statistical, or control treatments, the respondents are asked about their preferences on hypothetical policy packages generated by our randomized conjoint design. Each policy package is composed of its attributes $\boldsymbol{a} = (a_1, ..., a_m)$, where a_l, $l = 1, ..., m$, is a level of the lth attribute.

Consider a potential outcome experienced by respondent i given the information treatment status w_i, $Y_i\left(\boldsymbol{a}_j, \boldsymbol{a}_{-j}|w_i\right)$, where $\boldsymbol{a}_j$ and $\boldsymbol{a}_{-j}$ are alternative policy packages viewed by respondent i. $Y_i(\boldsymbol{a}_j, \boldsymbol{a}_{-j}|w_i)$ takes 1 if and only if a policy package $\boldsymbol{a}_j$ is preferred to an alternative package $\boldsymbol{a}_{-j}$, given the information treatment status w_i, such that

$$Y_i\left(\boldsymbol{a}_j, \boldsymbol{a}_{-j}|w_i\right) = \begin{cases} 1 & \text{if} \quad \boldsymbol{a}_j \succ_i \boldsymbol{a}_{-j} \\ 0 & \text{if} \quad \boldsymbol{a}_j \prec_i \boldsymbol{a}_{-j} \end{cases}$$

By randomly assigning the two information treatments and the control, we satisfy the unconfoundedness assumption,

$$w_i \perp\!\!\!\perp Y_i\left(\boldsymbol{a}_j, \boldsymbol{a}_{-j}|w_i\right) \big| \boldsymbol{X}_i,$$

where $\boldsymbol{X}_i$ denotes a vector of background characteristics of respondent i.

With an experiment with $Y_i = 1$, we observe $\boldsymbol{a}_j$, which was preferred, and with $Y_i = 0$, $\boldsymbol{a}_{-j}$, which was not preferred, conditional on $\boldsymbol{X}_i$ for each i. We treat $\boldsymbol{a}_j$ conditional on $\boldsymbol{X}_i$ and $\boldsymbol{a}_{-j}$ conditional on $\boldsymbol{X}_i$ as samples.

Then, let us define the individual information treatment effect for respondent i in period t as

$$\tau_i(\boldsymbol{a}_j, \boldsymbol{a}_{-j}) = Y_i(\boldsymbol{a}_j, \boldsymbol{a}_{-j}|w_i) - Y_i(\boldsymbol{a}_j, \boldsymbol{a}_{-j}|w_i^0),$$

where $w_i^0 \in \{c, n, s\}$ and $w_i^0 \neq w_i$, which captures the information treatment effects on respondent i's preferences over policy packages $\boldsymbol{a}_j$ and $\boldsymbol{a}_{-j}$. We thus consider the differences in policy preferences between the control condition ($w_i = c$), narrative-evidence treatment ($w_i = n$), and statistical-evidence treatment ($w_i = s$) as individual information treatment effect τ_i. Among them, we focus on the difference between the narrative-evidence and statistical-evidence treatments as the *narrative premium* of information treatment effects, denoted as follows,

$$\tau_i^{np}(\boldsymbol{a}_j, \boldsymbol{a}_{-j}) = Y_i(\boldsymbol{a}_j, \boldsymbol{a}_{-j}|w_i = n) - Y_i(\boldsymbol{a}_j, \boldsymbol{a}_{-j}|w_i = s). \tag{1}$$

4.1 Sorted group average treatment effects

Individual treatment effect τ_i is essentially a function of policy vector $\boldsymbol{a}$ and background characteristics vector $\boldsymbol{X}_i$ over $i = 1, \ldots, n$. Using information about respondent i's background characteristics $\boldsymbol{X}_i$ and respondent i's preferences over randomized policy attribute vectors $\boldsymbol{a}_j$ and $\boldsymbol{a}_{-j}$, we estimate the sorted group average marginal treatment effects (GATES) (Chernozhukov et al., 2018).

To estimate GATES, we classify samples into groups depending on predicted individual treatment effects. Note that a sample is a response of i, with background characteristics

$\boldsymbol{X}_i$, to policy package $\boldsymbol{a}_j$ or $\boldsymbol{a}_{-j}$. For respondent i, we observe $\boldsymbol{X}_i$ and responses to $\boldsymbol{a}_j$, which is the preferred policy package, and $\boldsymbol{X}_i$ and $\boldsymbol{a}_{-j}$, which is not preferred. As described in Section 3 above, our randomized conjoint experiment design imposes 5 tasks of choosing between two policy packages on each respondent. Thus we create 10 samples for each i.

Then, the GATES are defined as,

$$E\left[\tau_i(\boldsymbol{a}_j, \boldsymbol{a}_{-j})|G_k\right],$$

where $G_k = \{\tilde{\tau}_i \in [g_k, g_{k+1})\}$ is an indicator of a group membership, and $\tilde{\tau}_i$ is a predicted individual treatment effect conditional on background characteristics $\boldsymbol{X}_i$. We first estimate $\tilde{\tau}_i$ as a function of the vector of background characteristics $\boldsymbol{X}_i$ and policy attributes $\boldsymbol{a}_j$ and $\boldsymbol{a}_{-j}$, deploying the causal forest algorithm (Wager and Athey, 2018; Athey et al., 2019). We next sort out the samples to groups G_k for $k = 1, 2, 3$ by the degree of $\tilde{\tau}_i$ from the highest to lowest tertile. Here we sort samples to groups by $\tilde{\tau}_i$, not $|\tilde{\tau}_i|$ because information treatments may decrease preferences for a specific policy such that $\tilde{\tau}_i < 0$. Thus, our algorithm to estimate GATES is as follows:

1. Using the causal forest algorithm, predict individual information treatment effects on preferences for policy packages $\tilde{\tau}_i$ as a function of vector of background characteristics $\boldsymbol{X}_i$ and policy attributes $\boldsymbol{a}$.

2. Sort the samples into three groups, depending on policies' $\tilde{\tau}_i$. The highest tertile of $\tilde{\tau}_i$ is denoted Group 1 (G_1), and the lowest tertile is denoted Group 3 (G_3).

3. Estimate the GATES in each group.

Specifically, our estimate of interest is

$$E\left[\tau_i^{np}(\boldsymbol{a}_j, \boldsymbol{a}_{-j})|G_k^{np}\right], \tag{2}$$

where $\tilde{\tau}_i^{np}$ is the predicted value of the narrative premium, defined by Equation (1). $G_k^{np} = \{\tilde{\tau}_i^{np} \in [g_k^{np}, g_{k+1}^{np})\}$ indicates the tertile of $\tilde{\tau}_i^{np}$.

4.2 Classification analysis

GATES allow us to separate respondents by their responsiveness to narrative information. However, an additional step, classification analysis, is necessary to examine those who are more or less responsive. Specifically, comparing the background characteristics between those with high and low narrative premiums provides a good (albeit non-causal) summary relationship between individual traits and experimental outcomes.

To analyze the characteristics of groups, we estimate the following estimands,

$$E\left[Z|G_m^{np}\right] - E\left[Z|G_{\neg m}^{np}\right], \tag{3}$$

where G_m denotes one of the sorted groups from the GATES estimation, $G_{\neg m}$ denotes all or part of the other group(s), and Z denotes an element of the policy attribute vector $\boldsymbol{a}$ or background characteristics vector $\boldsymbol{X}_i$. For G_m and $G_{\neg m}$, we focus in practice on differences between G_1 and G_3, or the most and least affected groups.[10] By Equation (3), we subtract the expected value of Z among the samples classified in group $G_{\neg m}^{np}$ from that of those classified in group G_m^{np}. By doing so, we identify how much more likely Z is among the samples in group G_m^{np} than those in group $G_{\neg m}^{np}$, and hence, how much more likely Z is among the samples with highest narrative premium than those with least narrative premium than others.

[10]By "least affected", we mean the smallest value in an ordinal sense, as the minimum estimated treatment effect can be negative.

5 Results

5.1 Marginal Means

We first examine respondents' policy preferences with respect to government spending without any information treatment, i.e. in the control group. Figure 5 shows the marginal means, which is a descriptive measure that is identical to the conditional mean (Leeper et al., 2020). When attributes can take more than two levels or values, as is the case in our experimental design, the average marginal component effects (AMCE) estimate will vary depending on which value is used as the reference category. Marginal means, by contrast, show the average rating for each level, allowing us to draw inferences about the absolute level of each attribute-level's favorability.

Our analysis shows that for education, health care, and pensions, support linearly increases with spending. That is, respondents seem to believe that "the more the better". This suggests that there is broad consensus on expanding fiscal expenditures for public education, public health insurance, and the national pension plan. Meanwhile, support for poverty relief and infrastructure do not rise linearly. Notably, an increase from 15% to 20% lowered support for poverty relief at statistically significant levels. While the public health insurance and national pension plans cover all residents, poverty relief is direct income transfers for poor individuals only. Figure 5 thus supports our original intuition that enhanced spending for poverty relief programs are judged differently from other types of public goods.

Next, Figure 6 presents our main result regarding the narrative premium. Each panel shows the difference in marginal means between each information treatment and the control group, such that "narrative" indicates $w_i = n$ and "statistical" $w_i = s$ for $i = 1, 2, 3, \ldots$. It shows that the narrative evidence provided by information treatment 1 considerably

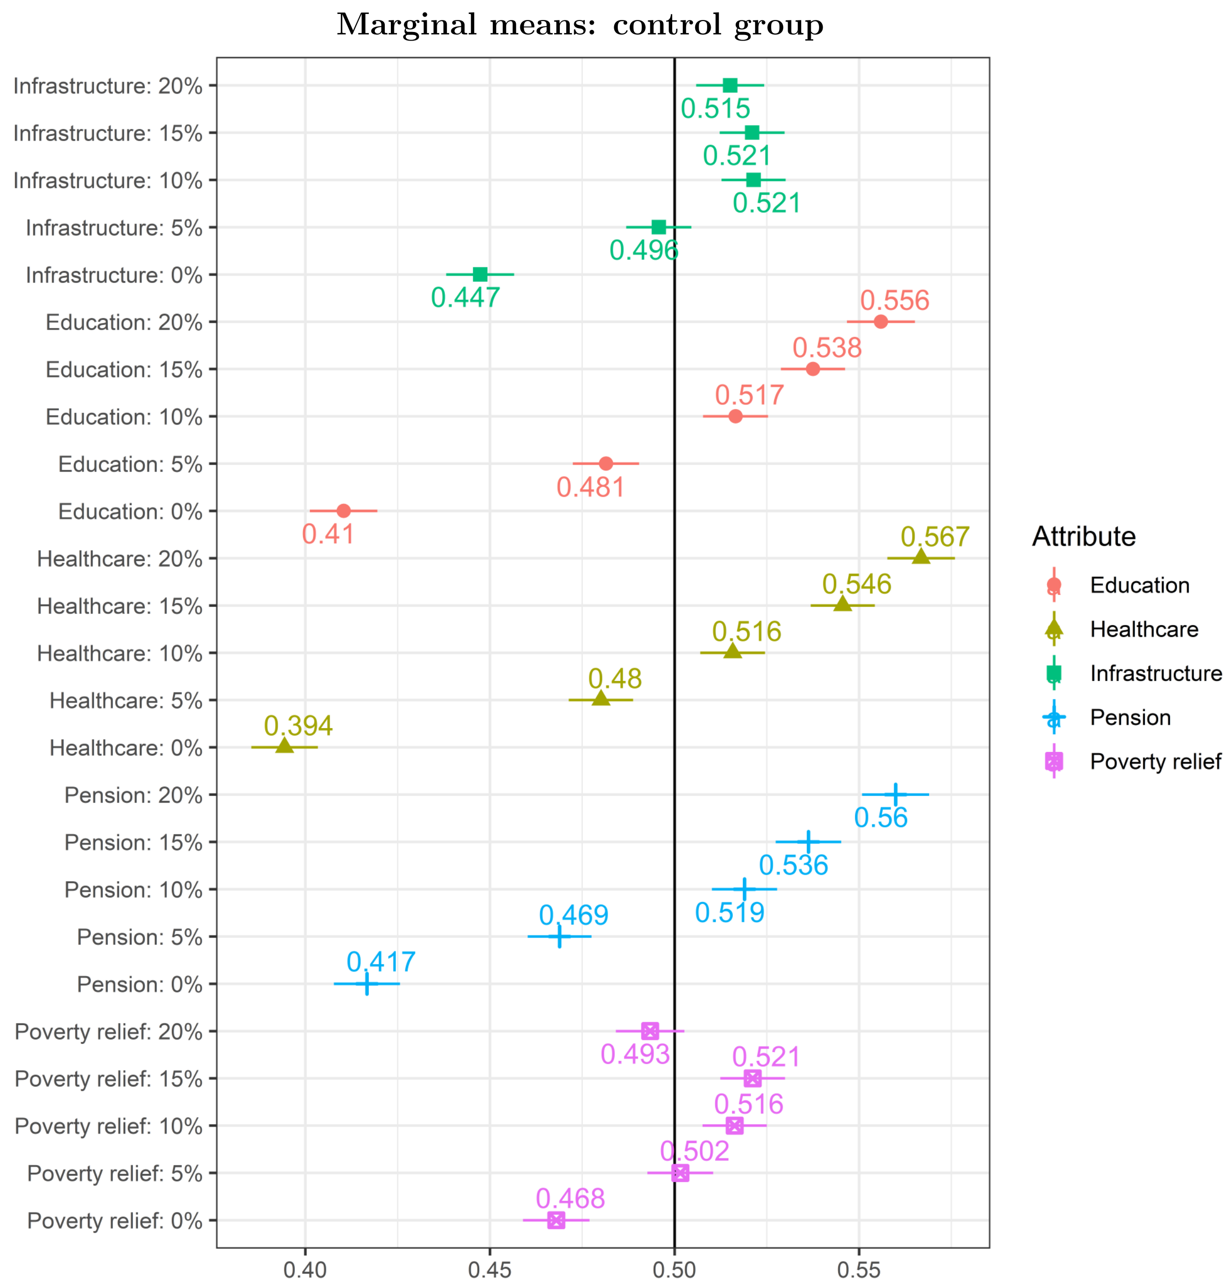

Figure 5: Estimates of marginal means of policy attributes in control group respondents. Estimates are based on OLS without a constant term and clustered standard errors. Bars denote 95% confidence intervals.

raises support for allocating 20% of additional tax revenues to poverty relief (0.028). While the statistical evidence from information treatment 2 had a positive impact for allocations of 20% (0.009), the impact of the narrative evidence was substantially larger. The gap between the two treatment effects captures the narrative premium defined by Equation (1).

Regarding other policies, information provision raised support for education. As with poverty relief, the impact of narrative evidence was greater than that of statistical evidence. This is likely related to the design of our narrative treatment. By highlighting the financial plight of single-parent households, respondents may have been persuaded to increase funding for education, which benefits the child. By contrast, we can see that the narrative evidence lowered support for high allocations (15% to 20%) to infrastructure. This indicates that respondents who preferred more support for poverty relief and education tended to finance the rise by reducing spending for infrastructure.

5.2 Subgroup Analyses: GATES and CLAN

This section explores heterogeneity in the magnitude of narrative premiums by reporting the estimated group sorted average treatment effects (GATES). Investigating treatment effect heterogeneity is difficult in conjoint experiments, as individuals rarely evaluate every possible attribute-level contribution. However, by leveraging recent advances in machine learning, we can apply the causal forest algorithm to observe heterogeneity in the effects of our information treatments on fiscal priorities.

Figure 7 presents the GATES of the narrative premium, defined in Equation (2) as "Narrative ($w_i = n$) vs. Statistical ($w_i = s$)" from the highest predicted tertile to the lowest tertile. This sorting is based on the difference in the predicted selection probability of a policy profile between those who receive the narrative versus statistical treatments, conditional on respondent characteristics and policy profile attributes. The depicted estimate is the size of the narrative premium, which we can see is significantly positive only

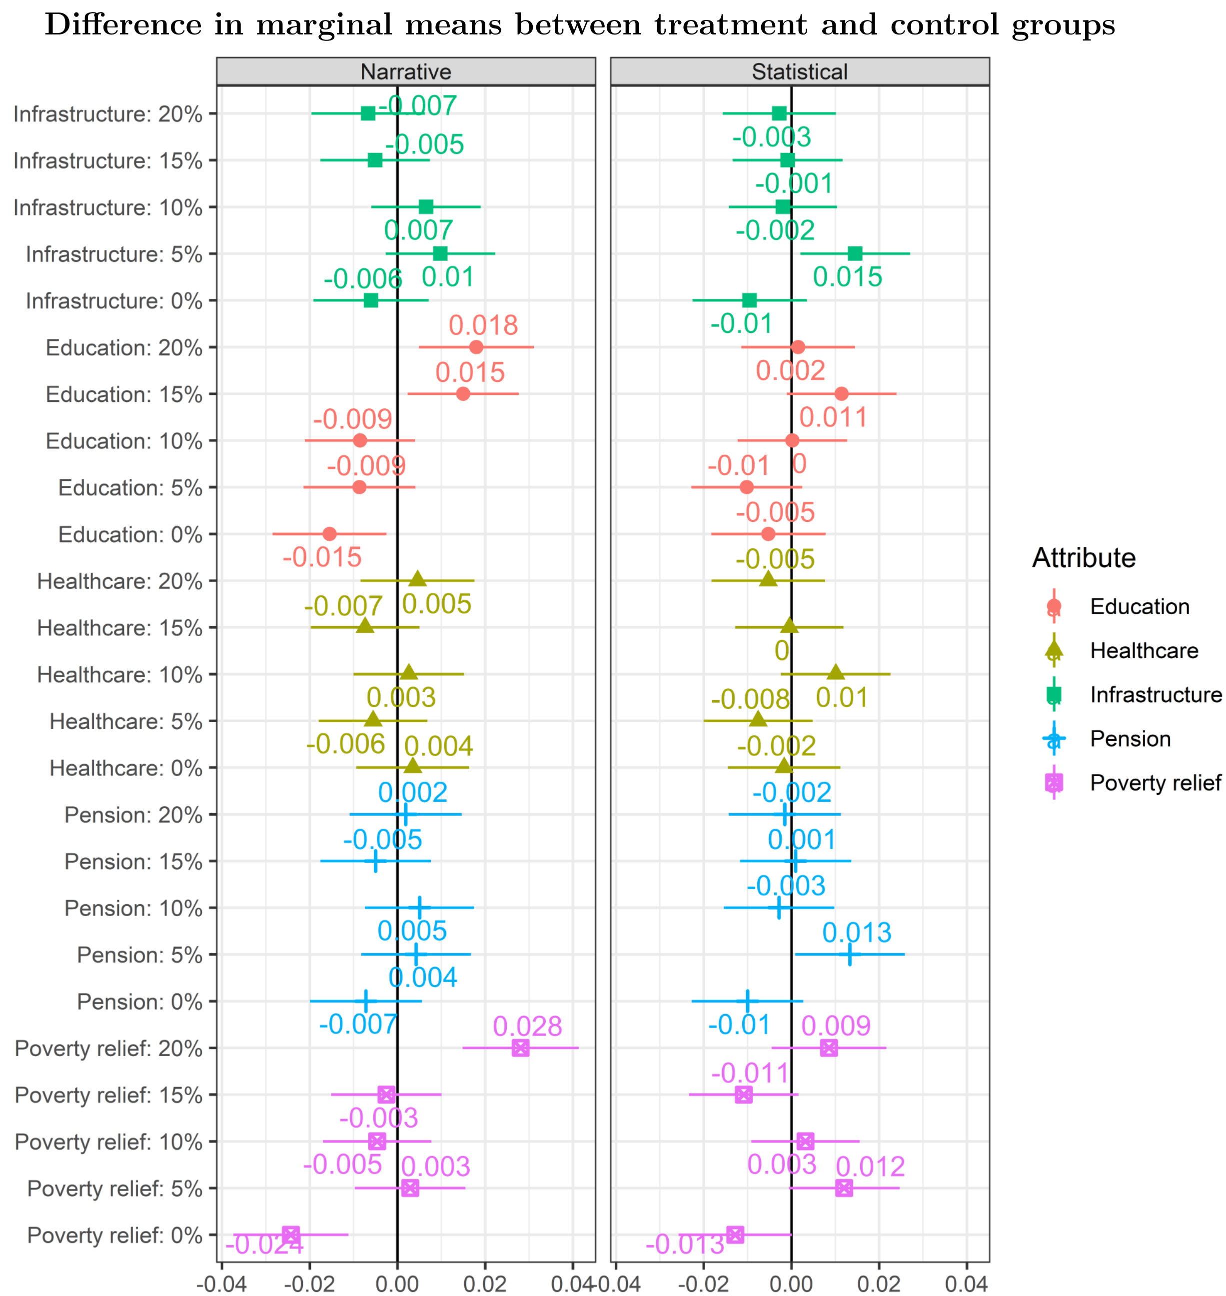

Figure 6: Estimates of the effects of the narrative (left panel) and statistical (right panel) treatments conditional on policy attribute levels. Estimates are based on differences-in-means between treatment and control groups. Bars denote 95% confidence intervals.

Group sorted average treatment effects of the narrative premium

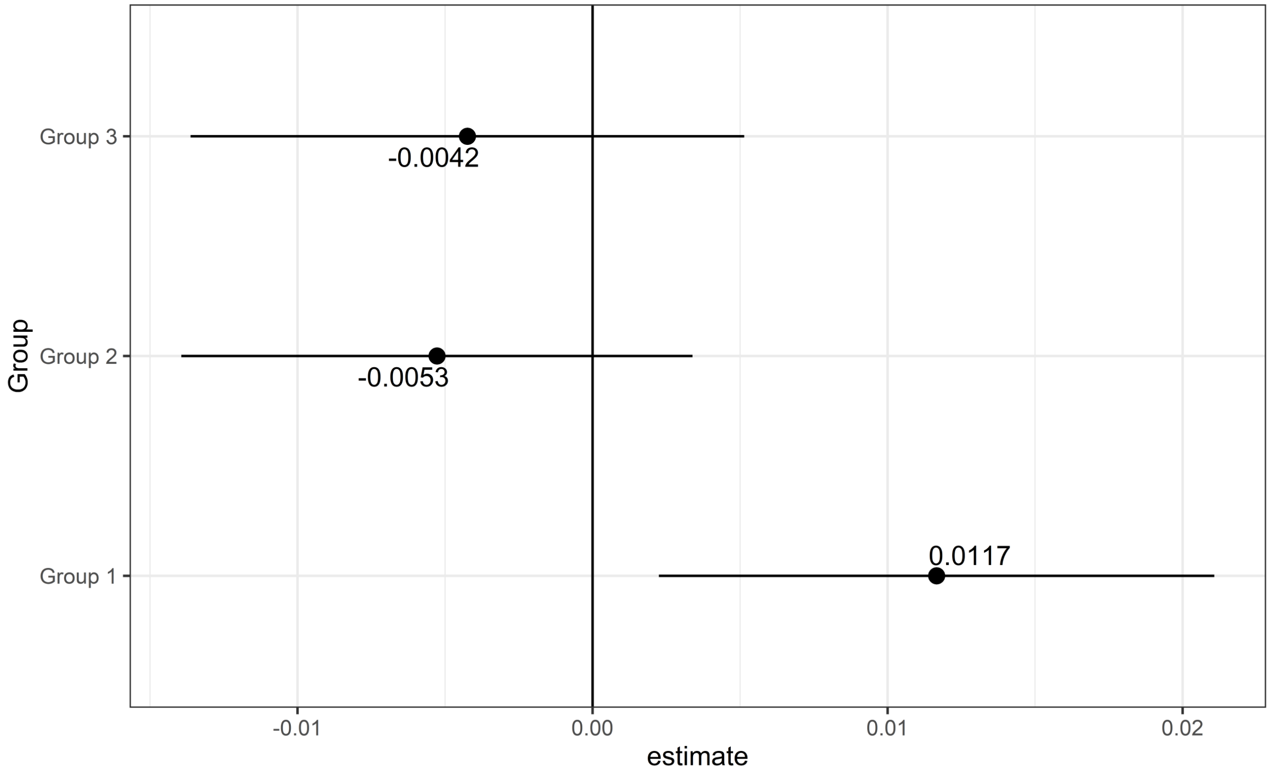

Figure 7: Estimates of the sorted group average affects of the narrative premium. These denote the group average of individual narrative effects on choice probabilities. Bars denote 95% confidence intervals. These groups are induced by the predicted individual narrative premium: Group 1 has the largest, Group 2 has intermediate, and Group 3 has the smallest predicted narrative effects.

for the highest predicted tertile, Group 1.

Using these GATES estimates, we can also examine differences between those whose narrative premium is greatest (Group 1) and least (Group 3) using classification analysis (CLAN). These factors can include differences in conjoint attribute levels, as well as characteristics of respondents. More specifically, based on the GATES, as a case of Equation (3), we first estimate the differences in means

$$E\left[Z|G_1^{np}\right] - E\left[Z|G_3^{np}\right], \tag{4}$$

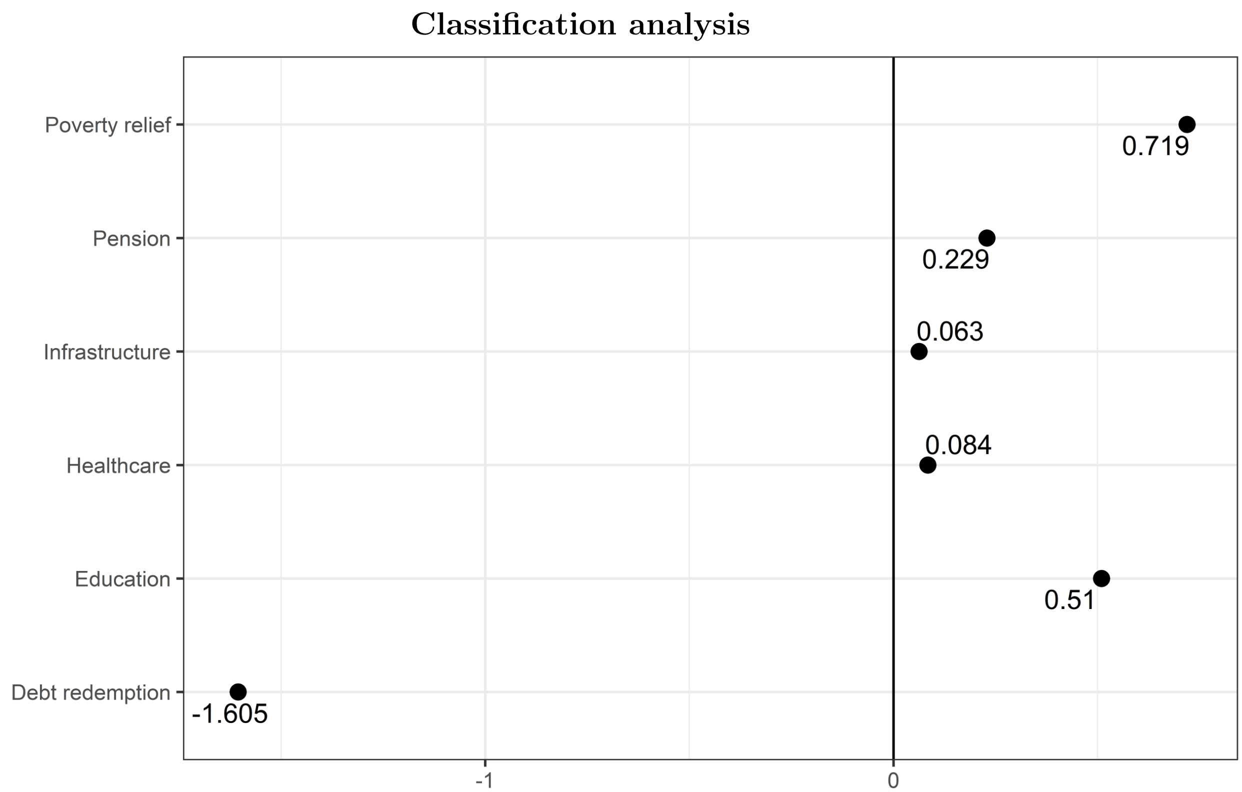

Figure 8: Results of classification analysis on policy attributes. The point estimates are the difference-in-means of attribute levels between the sorted Groups 1 and 3. Bars denote 95% confidence intervals.

as estimates of the following classification analyses such that Z ={support for spending for poverty relief programs}, {support for spending for national pension plan}, {support for infrastructure investment}, {support for national health insurance}, {support for spending for education}, or {support for debt redemption}. The estimates are displayed in Figure 8. Confidence intervals are technically shown, but due to their small size (related to our large sample size), they are obscured by the markers.

The horizontal axis of Figure 8 shows the value of Equation (4) for each Z. For example, the 0.719 estimate for poverty relief indicates that on average, those in Group 1 preferred policy packages with 0.719 percentage point (ppt) greater spending than those in Group

3. The positive estimates for education and modestly positive estimates for health and pension denote the same. By contrast, Group 1 respondents preferred packages with 1.605 ppt less money allocated to debt relief.

The results as a whole suggest that the narrative premium, or the increase in probability that a profile is selected when receiving the narrative over statistical treatment, served to boost support for greater spending on welfare programs for poor individuals, and to finance this by reducing repayment of government debt. The narrative premium also helped raise support for education expenditures, but it was relatively neutral for infrastructure investment. In sum, the narrative premium weakened preferences for fiscal entrenchment and strengthened spending for poverty relief and education.

Finally, we examine differences in respondent characteristics and attitudes between those with high and low narrative premiums. Figure 9 shows values of Equation (4) where $Z =$ {whether the respondent supports the LDP}, {respondent's perceived poverty rate}, {respondent's perceived child poverty rate},{respondent's number of children}, {respondent's individual income}, {respondent's household income}, {firm (employer) size where the respondent is employed}, {female dummy}, {respondent's educational background}, or {respondent's age}.

One key takeaway point is that the narrative premium exists for a broad cross-section of society and is not limited to specific political or socioeconomic groups. The narrative premium does not vary by partisan affinity in Japan, in contrast to highly partisan polities such as the United States. Similarly, the limited difference by income and education suggests that socioeconomic status is also less relevant. The positive value of age indicates that respondents with the highest narrative premium tended to be older than those who with the lowest.

There are, however, some notable exceptions. The positive values for expected poverty

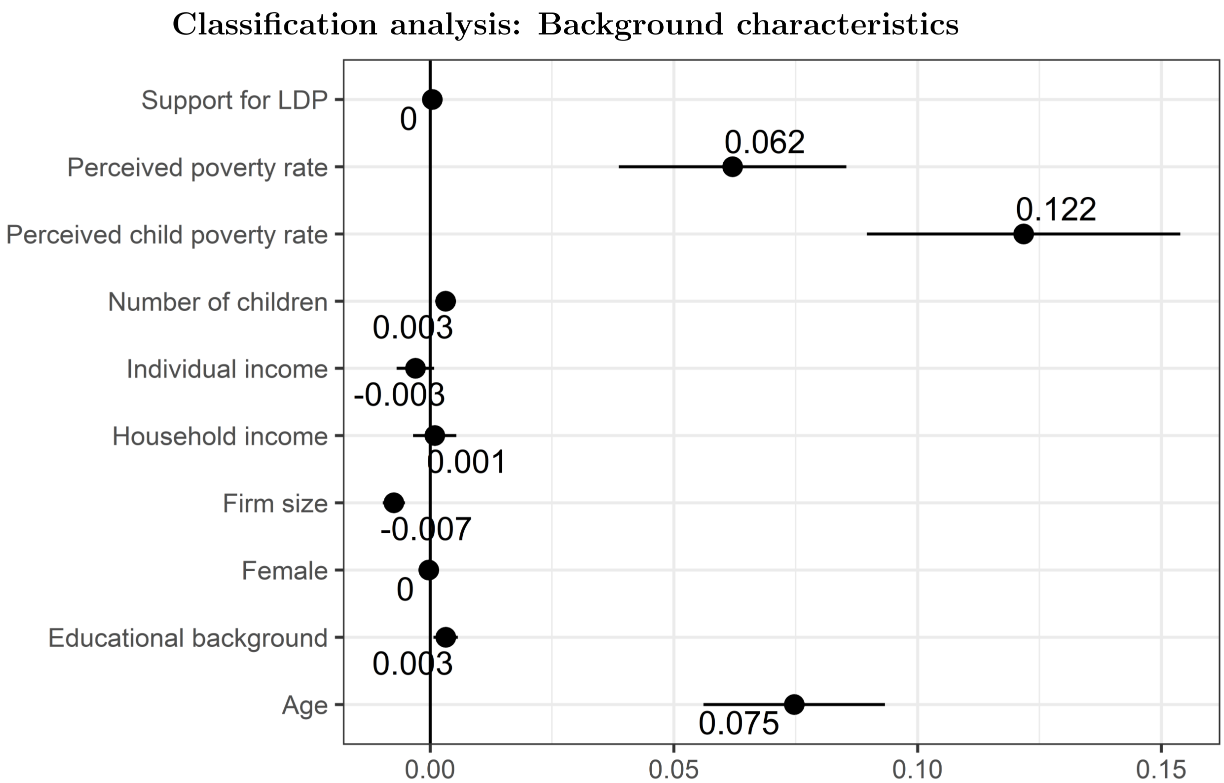

Figure 9: Results of the classification analysis on background characteristics. The point estimator is the difference-in-means of attributes levels between the sorted groups 1 and 3. Bars represent 95% confidence intervals. Categorical variables convert to continuous variables.

rate and expected child poverty rate mean that respondents with the highest narrative premium (Group 1) were also more likely to have expected higher rates on both. In other words, the narrative premium enhanced the poverty concerns of those who were already aware of it, rather than updating the knowledge of those who discounted prevailing poverty. This suggests that narrative information can convert the minds of those who are predisposed to oppose, or at least downgrade the priority of, poverty relief, in line with prior research (Slater and Rouner, 1996).

To be clear, this is not to say that these characteristics have no impact on the baseline support for those programs. Our results simply suggest that there are no estimated differences in the *size of the narrative premiums* by partisanship, income, or educational background.

6 Conclusion

We have demonstrated that the narrative premium, or the marginal effect of seeing narrative rather than statistical evidence, is substantial. Respondents are more likely to prefer higher levels of poverty relief spending when they are shown stories about the plight of single-parent families, rather than simply seeing data about their prevalence. These results have implications for our understanding of the politics of poverty relief and of political communication generally.

First, poverty relief is highly contentious in many countries. For one, it may be seen as disincentiving labor. For another, it intersects with prior stereotypes about the descriptive characteristics of the poor, such as "black welfare queens" in the United States (Gilens, 1999) or of immigrants who strategically move to countries with generous welfare provisions (Barrett and McCarthy, 2008; Fix and Haskins, 2002; Øland, 2019; Valentino et al., 2002). In this vein, a key concern is whether providing more "accurate" information about the nature of the poor—that they are primarily single-parent households—can sway public opinion.

Our analysis suggests that such information can be effective, particularly in a narrative form, because it highlights the *nature* of poverty, not because it alters prior beliefs about the *prevalence* of poverty. We find that the narrative premium is greater among those who overestimate actual levels of poverty. Put differently, what citizens may be less aware of is the manifold ways in which poverty influences the lives of destitute families.

This also connects to the external validity of our results. One possible cause of the narrative premium is prior perceptions of child poverty rates. Statistical evidence may be as—or perhaps more—influential if people tend towards underestimating real poverty rates. However, Figure 4 shows that on average, Japanese respondents' perceptions match objective trends, which stands in contrast to previous work based on Western cases. This suggests that the external validity of our findings may be greater in countries where social knowledge is high, but less where individuals lack accurate knowledge of income distributions.

Second, our study adds insights to the argument suggested by Slater and Rouner (1996) that narrative evidence is less likely to evoke psychological defense mechanisms. The consistent magnitude of narrative premiums across background characteristics, including partisanship, is consistent with this claim. In addition, our results indicate that narrative evidence is more likely to activate empathy among those who already have some knowledge about the issue. Our finding is similar to arguments that emotional reactions are independent from cognitive judgment, such as by Loewenstein et al. (2001), but different in a critical point. We find that narrative premium is not independent from prior perceptions, but rather increases with prior perception.

As a final point, we should note that our experimental design assesses just one facet of the determinants of public support for poverty relief. First, in order to add realism to the narrative and statistics treatments, we used a short essay for the former and a data graphic (pie chart) for the latter. Since news reports tend to use tables or figures when describing data, we chose the same for our statistics treatment. However, we cannot entirely discount the possibility that the medium—text versus pictures—influenced the results. Future studies may want to test whether the visual presentation style of information, not just its contents, alters treatment effects.

Second, our treatments differ slightly in how "households" are framed. The narrative essay discussed the effects of poverty on a single-parent mother and her child, while the statistical evidence showed the prevalence of poverty among single-parent households as collective units. It is possible that cues about individuals, especially children, are more evocative than those about abstract "households". A reasonable next step forward may be to design treatments that vary the makeup of households, such as those with and without children or with and without eldercare responsibilities. The demographic structure of families is diversifying in many countries, due to lengthening life expectancies and declining fertility rates. More work is necessary to disentangle how prior assumptions of and perceptions about what constitutes a "family" influences attitudes about providing additional welfare support.

References

Akerlof, George A. and Dennis J. Snower (2016) "Bread and bullets," *Journal of Economic Behavior & Organization*, 126, 58–71, https://doi.org/10.1016/j.jebo.2015.10.021, Thriving through Balance.

Alesina, Alberto, Stefanie Stantcheva, and Edoardo Teso (2018) "Intergenerational mobility and preferences for redistribution," *American Economic Review*, 108 (2), 521–54, https://doi.org/10.1257/aer.20162015.

Allen, Mike and Raymond W. Preiss (1997) "Comparing the persuasiveness of narrative and statistical evidence using meta-analysis," *Communication Research Reports*, 14 (2), 125–131, https://doi.org/10.1080/08824099709388654.

Athey, Susan and Guido W. Imbens (2019) "Machine learning methods that economists should know about," *Annual Review of Economics*, 11 (1), 685–725, https://doi.org/10.1146/annurev-economics-080217-053433.

Athey, Susan, Julie Tibshirani, and Stefan Wager (2019) "Generalized random forests," *Ann. Statist.*, 47 (2), 1148–1178, https://doi.org/10.1214/18-AOS1709.

Barrett, Alan and Yvonne McCarthy (2008) "Immigrants and welfare programmes: Exploring the interactions between immigrant characteristics, immigrant welfare dependence, and welfare policy," *Oxford Review of Economic Policy*, 24 (3), 542–559, https://doi.org/10.1093/oxrep/grn026.

Brader, Ted, Nicholas A. Valentino, and Elizabeth Suhay (2008) "What triggers public opposition to immigration? Anxiety, group cues, and immigration threat," *American Journal of Political Science*, 52 (4), 959–978, https://doi.org/10.1111/j.1540-5907.2008.00353.x.

Chater, Nick and George Loewenstein (2016) "The under-appreciated drive for sense-making," *Journal of Economic Behavior & Organization*, 126, 137–154, https://doi.org/10.1016/j.jebo.2015.10.016, Thriving through Balance.

Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Iván Fernández-Val (2018) "Generic machine learning inference on heterogeneous treatment effects in randomized experiments," Working Paper 24678, National Bureau of Economic Research, https://doi.org/10.3386/w24678.

Chiavacci, David (2008) "From class struggle to general middle-class society to divided society: Societal models of inequality in postwar Japan," *Social Science Japan Journal*, 11 (1), 5–27, https://doi.org/10.1093/ssjj/jyn022.

Chudy, Jennifer (2021) "Racial sympathy and its political consequences," *The Journal of Politics*, 83 (1), 122–136, https://doi.org/10.1086/708953.

Clark, Andrew and Conchita D'Ambrosio (2015) "Attitudes to income inequality: Experimental and survey evidence," *Handbook of Income Distribution*, 2, 1147–1208, https://doi.org/10.1016/B978-0-444-59428-0.00014-X.

De Wit, John B. F., Enny Das, and Raymond Vet (2008) "What works best: Objective statistics or a personal testimonial? An assessment of the persuasive effects of different types of message evidence on risk perception," *Health Psychology*, 27 (1), 110–115, https://doi/10.1037/0278-6133.27.1.110.

Diaz-Cayeros, Alberto, Federico Estévez, and Beatriz Magaloni (2016) *The Political Logic of Poverty Relief: Electoral Strategies and Social Policy in Mexico*, New York: Cambridge University Press.

Esarey, Justin, Tim Salmon, and Charles Barrilleaux (2012) "Social insurance and income redistribution in a laboratory experiment," *Political Research Quarterly*, 65 (3), 685–698, https://doi.org/10.1177/1065912911411096.

Estevez-Abe, Margarita (2008) *Welfare and Capitalism in Postwar Japan*, Cambridge: Cambridge University Press.

Ferraro, Paul J., Juan Jose Miranda, and Michael K. Price (2011) "The persistence of treatment effects with norm-based policy instruments: Evidence from a randomized environmental policy experiment," *American Economic Review*, 101 (3), 318–322, https://doi.org/10.1257/aer.101.3.318.

Ferraro, Paul J. and Michael K. Price (2013) "Using nonpecuniary strategies to influence behavior: Evidence from a large-scale field experiment," *The Review of Economics & Statistics*, 95 (1), 64–73, https://doi.org/10.1162/REST_a_00344.

Fix, Michael and Ron Haskins (2002) "Welfare benefits for non-citizens,"Technical report, Brookings, `https://www.brookings.edu/research/welfare-benefits-for-non-citizens/`.

Gilens, Martin (1996) ""Race coding" and white opposition to welfare," *The American Political Science Review*, 90 (3), 593–604, https://doi.org/10.2307/2082611.

——— (1999) *Why Americans Hate Welfare: Race, Media, and the Politics of Antipoverty Policy*, Chicago: University of Chicago Press, `https://press.uchicago.edu/ucp/books/book/chicago/W/bo3633527.html`.

Guo, Yaqin, Peta Ashworth, Yan Sun, Bo Yang, Jinhua Yang, and Jiaowei Chen (2019) "The influence of narrative versus statistical evidence on public perception towards CCS in China: Survey results from local residents in Shandong and Henan provinces," *International Journal of Greenhouse Gas Control*, 84, 54–61, https://doi.org/10.1016/j.ijggc.2019.02.021.

Hommerich, Carola and Toru Kikkawa (2019) "Movement behind the scenes: The quiet transformation of status identification in Japan," *Social Science Japan Journal*, 22 (1), 11–24, https://doi.org/10.1093/ssjj/jyy041.

Hong, Seoyeon and Hee Sun Park (2012) "Computer-mediated persuasion in online reviews: Statistical versus narrative evidence," *Computers in Human Behavior*, 28 (3), 906–919, https://doi.org/10.1016/j.chb.2011.12.011.

Hopkins, Daniel J., John Sides, and Jack Citrin (2019) "The Muted Consequences of Correct Information about Immigration," *The Journal of Politics*, 81 (1), 315–320, 10.1086/699914.

Imbens, Guido W. and Donald B. Rubin (2015) *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*, Cambridge: Cambridge University Press.

Inaba, Miyuki (2011) "Increasing Poverty in Japan: Social Policy and Public Assistance Program," *Asian Social Work and Policy Review*, 5 (2), 79–91, https://doi.org/10.1111/j.1753-1411.2011.00050.x.

Jusko, Karen Long (2017) *Who Speaks for the Poor? Electoral Geography, Party Entry, and Representation*, New York: Cambridge University Press.

Kalla, Joshua L. and David E. Broockman (2020) "Reducing exclusionary attitudes through interpersonal conversation: Evidence from three field experiments," *American Political Science Review*, 114 (2), 410–425, https://doi.org/10.1017/S0003055419000923.

Kanbayashi, Hiroshi (2019) "The changing images of Japan's social stratification: The other side of the 'quiet transformation'," *Social Science Japan Journal*, 22 (1), 45–63, https://doi.org/10.1093/ssjj/jyy048.

Kerschbamer, Rudolf and Daniel Müller (2020) "Social preferences and political attitudes: An online experiment on a large heterogeneous sample," *Journal of Public Economics*, 182, 104076, https://doi.org/10.1016/j.jpubeco.2019.104076.

Kohara, Miki, Midori Matsushima, and Fumio Ohtake (2019) "Effect of unemployment on infant health," *Journal of the Japanese and International Economies*, 52, 68–77, https://doi.org/10.1016/j.jjie.2019.03.002.

Krueger, Alan B. and Bruce D. Meyer (2002) "Chapter 33 Labor supply effects of social insurance," in *Handbook of Public Economics*, 4, 2327–2392: Elsevier, https://doi.org/10.1016/S1573-4420(02)80012-X.

Kuziemko, Ilyana, Michael I. Norton, Emmanuel Saez, and Stefanie Stantcheva (2015) "How elastic are preferences for redistribution? Evidence from randomized survey experiments," *American Economic Review*, 105 (4), 1478–1508, https://doi.org/10.1257/aer.20130360.

Leeper, Thomas J., Sara B. Hobolt, and James Tilley (2020) "Measuring subgroup preferences in conjoint experiments," *Political Analysis*, 28 (2), 207–221, https://doi.org/10.1017/pan.2019.30.

Loewenstein, George F., Elke U. Weber, Christopher K. Hsee, and Ned Welch (2001) "Risk as feelings," *Psychological Bulletin*, 127 (2), 267–286, https://doi.org/10.1037/0033-2909.127.2.267.

Maldonado, Laurie C. and Rense Nieuwenhuis (2015) "Family policies and single parent poverty in 18 OECD countries, 1978–2008," *Community, Work & Family*, 18 (4), 395–415, 10.1080/13668803.2015.1080661.

Mares, Isabela (2003) *The Politics of Social Risk: Business and Welfare State Development*, Cambridge Studies in Comparative Politics, Cambridge: Cambridge University Press.

Mazor, Kathleen M., Joann Baril, Elizabeth Dugan, Frederick Spencer, Pamela Burgwinkle, and Jerry H. Gurwitz (2007) "Patient education about anticoagulant medication: Is narrative evidence or statistical evidence more effective?," *Patient Education and Counseling*, 69 (1), 145–157, https://doi.org/10.1016/j.pec.2007.08.010.

Moriguchi, Chiaki and Emmanuel Saez (2008) "The evolution of income concentration in Japan, 1886–2005: Evidence from income tax statistics," *The Review of Economics & Statistics*, 90 (4), 713–734, https://doi.org/10.1162/rest.90.4.713.

Nakamura, Sayaka (2014) "Parental income and child health in Japan," *Journal of the Japanese and International Economies*, 32, 42–55, https://doi.org/10.1016/j.jjie.2013.12.003.

Nan, Xiaoli, Michael F. Dahlstrom, Adam Richards, and Sarani Rangarajan (2015) "Influence of evidence type and narrative type on HPV risk perception and intention to obtain the HPV vaccine," *Health Communication*, 30 (3), 301–308, https://doi.org/10.1080/10410236.2014.888629, PMID: 25061715.

Nyhan, Brendan (2020) "Facts and Myths about Misperceptions," *The Journal of Economic Perspectives*, 34 (3), 220–236, `https://www.jstor.org/stable/26923548`.

Øland, Trine (2019) *Welfare Work with Immigrants and Refugees in a Social Democratic Welfare State*, London: Routledge, https://doi.org/10.4324/9781351264440.

Shiller, Robert J. (2017) "Narrative economics," *The American Economic Reivew*, 107 (4), 967–1004, http://dx.doi.org/10.1257/aer.107.4.967.

Shirahase, Sawako (2014) *Social Inequality in Japan*, London: Routledge.

——— (2021) "Social Stratification Theory and Population Aging Reconsidered," *Social Science Japan Journal*, 24 (2), 277–288, 10.1093/ssjj/jyab010.

Slater, Michael D. and Donna Rouner (1996) "Value-affirmative and value-protective processing of alcohol education messages that include statistical evidence or anecdotes," *Communication Research*, 23 (2), 210–235, https://doi.org/10.1177/009365096023002003.

Tachibanaki, Toshiki (2006) "Inequality and poverty in Japan," *The Japanese Economic Review*, 57 (1), 1–27, https://doi.org/10.1111/j.1468-5876.2006.00355.x.

Tanaka, Masako (2022) *Zōzei no Gōi-Keisei: Renritsu Seiken Jidai no Seitō-kan Kyōsō to Kyōchō*, Tokyo: Nippon Hyōron-sha.

Tatsiramos, Konstantinos and Jan C. van Ours (2014) "Labor market effects of unemployment insurance design," *Journal of Economic Surveys*, 28 (2), 284–311, https://doi.org/10.1111/joes.12005.

Tversky, Amos and Daniel Kahneman (1973) "Availability: A heuristic for judging frequency and probability," *Cognitive Psychology*, 5 (2), 207–232, https://doi.org/10.1016/0010-0285(73)90033-9.

Valentino, Nicholas A., Vincent L. Hutchings, and Ismail K. White (2002) "Cues that Matter: How Political Ads Prime Racial Attitudes During Campaigns," *American Political Science Review*, 96 (1), 75–90, https://doi.org/10.1017/S0003055402004240.

Wager, Stefan and Susan Athey (2018) "Estimation and inference of heterogeneous treatment effects using random forests," *Journal of the American Statistical Association*, 113 (523), 1228–1242, https://doi.org/10.1080/01621459.2017.1319839.

Winterbottom, Anna, Hilary L. Bekker, Mark Conner, and Andrew Mooney (2008) "Does narrative information bias individual's decision making? A systematic review," *Social Science & Medicine*, 67 (12), 2079 – 2088, https://doi.org/10.1016/j.socscimed.2008.09.037.

Yuasa, Makoto (2017) *"Nantoka Suru" Kodomo no Hinkon ("Do Something" with Child Poverty)*, Tokyo: Kadokawa.

Online Appendix A

Table A1 shows the household income distribution in a survey administered by the Ministry of Health, Labour and Welfare in 2019.

Table A1: Distribution of household income in the National Livelihood Survey

Income level	N	Share
Total	10,000	100.00%
Less than 0.50 million yen	120	1.20%
0.50–0.99 million yen	519	5.19%
1.01–1.49 million yen	631	6.31%
1.50–1.99 million yen	632	6.32%
2.00–2.49 million yen	689	6.89%
2.50–2.99 million yen	666	6.66%
3.00–3.49 million yen	711	7.11%
3.50–3.99 million yen	574	5.74%
4.00–4.49 million yen	555	5.55%
4.50–4.99 million yen	491	4.91%
5.00–5.49 million yen	488	4.88%
5.50–5.99 million yen	380	3.80%
6.00–6.49 million yen	463	4.63%
6.50–6.99 million yen	344	3.44%
7.00–7.49 million yen	329	3.29%
7.50–7.99 million yen	288	2.88%
8.00–8.49 million yen	260	2.60%
8.50–8.99 million yen	232	2.32%
9.00–9.49 million yen	216	2.16%
9.50–9.99 million yen	185	1.85%
10.00 million or over	1,225	12.25%

Source: National Livelihood Survey 2019 by the Ministry of Health, Labour and Welfare, Japan `https://www.e-stat.go.jp/stat-search/file-download?statInfId=000031957851&fileKind=1` (Accessed on September 10, 2021).

Online Appendix B: Research Ethics

Our survey experiment was conducted between February 28 and March 14, 2020. Respondents were recruited from the national panel of Rakuten Insight, a major survey vendor in Japan with 2.2 million registrants. We did not stratify our sample by demographic characteristics, but given our large sample size (15,000), our respondent pool did not differ greatly from census distributions on gender or age, as described in Table 3 in the main text.

Detailed information about Rakuten Insight's respondent pool is available from its website. https://insight.rakuten.co.jp/download/PanelProfile_EN.pdf and https://insight.rakuten.co.jp/download/PanelCharacteristicSurveyEN.pdf.

When using Internet survey samples, we must be careful of biases in respondent demographics. Nagayoshi et al.'s (2020) comparison of online surveys using Rakuten Insight and a mail survey using random sampling finds no significant differences by demographic variables or social awareness. In addition, many prior studies in political science have used quota sampling from Rakuten Insight (e.g. Igarashi et al. 2022). While we cannot fully discount other underlying biases inherent to internet survey populations, we do not believe that Rakuten Insight's respondent pool deviates greatly from other survey providers.

This survey experiment was approved by the Institutional Review Board of ANONYMIZED (Approval number: ANONYMIZED). The survey was funded by MEXT/JSPS KAKENHI ANONYMIZED. A pre-analysis plan was not registered, although information about the purpose and design of the study was provided to the IRB. Personal information that could identify respondents, beyond basic demographic characteristics, were not collected in the survey or shared by Rakuten Insight. Consent was obtained before starting the survey, and the debriefing page explained the purpose and structure of the survey. Respondents were not presented with any false or deceptive information.

All participants who completed the survey received Rakuten points that could be used on Rakuten Ichiba, a major online shopping portal. The financial value of these points is proprietary information of Rakuten, and was not disclosed to the PIs.

Igarashi, Akira, Hirofumi Miwa, and Yoshikuni Ono. 2022. "Why Do Citizens Prefer High-Skilled Immigrants to Low-Skilled Immigrants? Identifying Causal Mechanisms of Immigration Preferences with a Survey Experiment." *Research & Politics*, Online First.

Nagayoshi, Kikuko, Mitsuru Matsutani, and Naoto Higuchi. 2020. "An Online Survey with a Large Sample-Size: A case of the survey for demonstrators after the Great East Japan Earthquake", *Political Studies* 35(1): 145-158.