

Can Incorrect Beliefs about the Racial Composition of Welfare and Unemployment Insurance Beneficiaries be Changed?

By JEFFREY CARPENTER, JAKINA DEBNAM GUZMAN, PETER HANS MATTHEWS, AND ERIN L. WOLCOTT*

Some argue that support for the social safety net in the United States is influenced by beliefs about the beneficiaries' race. Information treatments have the potential to change these beliefs, but for them to be policy relevant, their effects must last beyond the intervention. Our findings from two parallel experiments that exploit the different racialized histories of welfare and unemployment insurance indicate that racial beliefs do predict stated support for the racially stigmatized welfare program but not for the less stigmatized unemployment program. We also find these beliefs are stable if uncorrected and that they can be persistently corrected.

JEL: J08, J15, D9

Keywords: Experiment, Information Provision, Welfare, Unemployment Insurance, Social Safety Net, Discrimination

Since Quadagno (1994) and Gilens (1995), it has become commonplace in the social sciences to understand declining support for social safety net policies in the United States as a partial consequence of Americans' racialized beliefs about the beneficiaries of welfare and other programs. McGhee (2022) and Wolcott (2024), for example, argue that the decline was a backlash to the increased access of Blacks to various programs following the Civil Rights Movement. In this paper, we provide a partial report on a large field experiment designed to generate causal test(s) of this provocative proposition. Our immediate focus is the explanatory power of beliefs about the racial composition of welfare (TANF) and unemployment insurance (UI) beneficiaries and their *persistent* manipulation.¹ In later work (Carpenter et al., 2024a), we show that the persistent information

* Carpenter: Middlebury College and IZA, jcarpent@middlebury.edu. Guzman: Amherst College, jguzman@amherst.edu. Matthews: Middlebury College and Aalto University, pmatthew@middlebury.edu. Wolcott: Middlebury College, wolcott@middlebury.edu. For excellent research assistance, we thank Allison Lounsbury, Blair Jai and Bishal Panthi. For helpful comments, we are grateful to Dania Francis, seminar participants at Williams College, and to participants at the National Economic Association session of the 2025 AEA Annual Meetings and the Sixth World Conference on Remedies to Racial and Social Inequality. The trial was pre-registered with the Social Science Registry under ID AEARCTR-0013774 (Carpenter, Debnam Guzman, Matthews, and Wolcott 2024) and approved by the Middlebury IRB (Protocol #125).

¹Authors studying the U.S. have used “welfare” to describe programs including Aid to Families with Dependent Children (AFDC), Temporary Assistance to Needy Families (TANF), Medicaid, the Supplemental Nutrition Assistance Program (SNAP), Housing Assistance, and Supplemental Security Income (SSI). We focus on the program most commonly associated with welfare so as not to confuse our participants.

effects we find shift policy preferences. Our results therefore challenge the now conventional wisdom that such preferences are inelastic (Kuziemko et al., 2015).

We focus on unemployment insurance and welfare because their racialized histories diverged early on. Both programs were established by the Social Security Act of 1935 and left state governments to set benefit amounts and eligibility requirements. Starting in 1939, Aid to Families with Dependent Children (the program that became TANF) disproportionately served Black families because widows and children of retired non-farm workers—who were mostly White—became eligible for old-age assistance (Quadagno, 1994). Our expectation is that beliefs about the proportion of program beneficiaries who are Black will matter more for support of welfare than they do for support of unemployment benefits because of a racialized conception of “deservingness,” in which UI recipients are often viewed as the working victims of “bad luck”, while TANF recipients are perceived as doing little to “earn” their transfers, notwithstanding some work requirements.²

In what follows, we first show that beliefs about the number of welfare recipients who are Black is a robust predictor of how much respondents think TANF benefits should be increased or decreased. Based on previous work (Robbett and Matthews, 2018), we are confident that our incentives attenuated the influence of “expressive beliefs.” Consistent with our pre-analysis plan, we also perform a similar exercise for unemployment benefits, and show that no such correlation exists, which we view as *prima facie* evidence of policy-specific differences in the conception of deservingness, and its interaction with race (Fong, 2007). Second, we show that these beliefs are stable from month to month, in the sense that untreated beliefs also predict incentivized beliefs elicited one month later. This means that they have the potential to support policy preferences. At the same time, and critical for the broader project, we show that we can induce persistent changes in beliefs about the composition of TANF and UI recipients. To do so, we adopt an information provision protocol (*e.g.*, Haaland, Roth and Wohlfart, 2023) that exploits the substantial variation in misbeliefs and randomly corrects these priors for half the population. If treated participants fully updated their posteriors, beliefs in the two rounds should be uncorrelated, and this is precisely what we find.

We are not the first, of course, to use an information provision protocol to explore beliefs about race and social safety net policy preferences. Alesina, Ferroni and Stantcheva (2021), for example, find that information about systemic racism, but not Black-White earnings gaps, affects policy support. Using a protocol similar to ours, Akesson et al. (2022) find that priming White respondents to think about the racial composition of welfare reduces program support, but that the provision of accurate information has no effect. Some studies have also explored the persistence of racial beliefs: Callaghan et al. (2021), for example, find that participants updated their beliefs about the White-Black wealth gap, and that

²Carpenter et al. (2024a) indeed find that, on average, experimental participants state that beneficiaries of UI are more deserving of help than are beneficiaries of TANF.

the information was retained for at least 18 months. Likewise, Haaland and Roth (2023) find that respondents retain new information about racial discrimination. The important differences between our work and these antecedents are that we induce persistent changes in beliefs, we document differences in the relationship between beliefs across policies and, in related work (Carpenter et al., 2024a), we demonstrate that the changes in beliefs move policy preferences, confirming the existence of a causal link between beliefs and preferences.

I. Experimental Design

This section provides a high-level overview of our experimental protocol described in our pre-analysis plan (Carpenter et al., 2024b): for complete details, see Carpenter et al. (2024a). During the summer of 2024, we designed and implemented two parallel three-stage information provision experiments on Connect³ targeting a nationally representative sample of participants in the United States. In the experiments, participants were asked for their beliefs about how many Black people made use of either TANF or UI in 2021. The participants were then asked to report their support for the social safety net program on which their experiment focused. Prior to reporting their support, half the participants were chosen at random to be told the correct number.

Stage 1 (which ran during the second week of June) collected individual characteristics and asked participants about their implicit and explicit racial preferences using the familiar Implicit Association Test and a likert-response question. The median Stage 1 participant spent 6.9 minutes and earned a flat fee of \$1.25 (the equivalent of \$10.70 per hour).

To minimize any priming or experimenter demand effects, we started Stage 2 one month after Stage 1 was completed. Without any cajoling, 94% of the Stage 1 respondents returned for Stage 2 and were randomly sorted into one of the four experimental conditions.⁴ Stage 2 began by informing all participants that they would be paid a \$1 bonus if they correctly (i.e., true number ± 2) responded to one of the following two belief questions to which they were randomized: (1) *Out of every 100 adults who received welfare from the U.S. government (sometimes referred to as TANF or Temporary Assistance for Needy Families) in 2021, how many do you think identified as Black?* or (2) *Out of every 100 adults who received unemployment benefits from the U.S. government in 2021, how many do you think identified as Black?* In addition, to attenuate base rate neglect, all participants were informed that during 2021, 13 out of every 100 adults in the U.S. identified as Black. After answering, participants revealed (on a scale from 0 to 10) how confident they were in their belief.

Those participants randomized into the treatment conditions were then told the correct information (29 people for TANF and 18 for UI). Directly after stating

³Connect is a crowdsourcing platform run by CloudResearch which provides participants for online surveys.

⁴As expected with 94% returning, there is no selection on observables into Stage 2.

their beliefs and having them corrected or confirmed in the treatment conditions, we asked participants our policy support questions. The questions told participants what the typical TANF or UI benefit was per month in 2021 and asked them how the benefit should change.⁵ They could respond on a vertical slider anywhere from -100% (end the benefit altogether) to $+100\%$ (double the size of the benefit). The median time to completion in Stage 2 was 4.1 minutes, for which participants received a flat fee of \$1 and the potential to earn up to an additional \$2 in bonuses (the equivalent of up to \$43.90 per hour).

Stage 3 of the experiment, run a month after Stage 2, was designed to test the persistence of our information interventions. Participants were given a \$1 flat fee and provided additional incentives (like those in Stage 2) to provide their current beliefs about the number of Black people who used the programs in 2021. Importantly, all bonus payments were made after Stage 3 was completed so that participants in the control groups were not inadvertently informed about the actual rates of Black participants in the two social safety net programs.

With the exception of the number of participants aged 65+, our respondents are representative of the U.S. population. Further, as detailed in Carpenter et al. (2024a), we achieved treatment balance in both experiments. Information about sample size and power calculations may be found in Online Appendix Table A1 and the pre-analysis plan.

II. Results

A. Beliefs

Participants' prior beliefs from the start of Stage 2 about the percent of TANF and UI beneficiaries who were Black vary from 1 to 100 with a mean of 30.08, which is only significantly different from the true number (29) at the 5% level ($t = 2.03$, $p = 0.04$). Put differently, the beliefs about TANF participation of our participants are very close to accurate, on average. Importantly, from an experimental design and inference perspective, 54% of welfare beliefs are below the correct number and 46% are at or above the correct number.

Considering the determinants of participant TANF beliefs, Online Appendix Table A2 suggests that there is no significant correlation between the level of one's belief and one's confidence in that belief, though the absolute difference between one's belief and the true number is positively correlated with confidence at the 10% level. In terms of demographics, we find that older participants have lower and less accurate beliefs, while women and participants with an implicit

⁵According to the Center on Budget and Policy Priorities, the maximum TANF benefit paid to a mother (with two children) in the median state was \$498 a month at the end of 2021. According to the U.S. Department of Labor, the maximum UI benefit paid to a claimant in the median state was \$1852 a month at the end of 2021. We chose 2021 because, at the time of the experiment, it is the most recent year the Census Bureau reported the racial composition of TANF and UI recipients at www.census.gov/library/visualizations/interactive/social-safety-net-benefits.html.

bias against Black people have larger misbeliefs.⁶

Beliefs about the number of Black UI participants vary from 0 to 98 with a mean of 22.95, which is significantly higher than the true value of 18 ($t = 10.22$, $p < 0.01$). That said, we see that the beliefs are balanced almost equally on either side of the actual number (the median UI belief is 17), so there were as many participants receiving news that their beliefs were too high as received the opposite news. UI belief confidence averages 4.33 (on a scale from 0-10 where 10 is “certain”) and is a poor predictor of accuracy, as seen in Table A2. Table A2 also indicates that UI beliefs are lower for older and White respondents after controlling for participant explicit preference for Whites, which are a strong predictor of UI beliefs. Concerning accuracy, older participants hold more accurate priors, as do people with a college degree, while women and explicit racists tend to have less accurate beliefs.

B. Policy support

In the welfare experiment, participants wish to increase TANF benefits by a third (34%), on average. Fewer people would like to dramatically increase unemployment benefits. Here, the modal response is to increase UI by about a quarter and the average suggested increase is 22%. The left panel of of Figure 1 depicts a strong association between participant misbeliefs about the racial composition of welfare participation and support for the program. Specifically, participants who believe that there are considerably fewer Black people collecting welfare than is correct tend to want to increase benefits by as much as 50%, while those who overestimate the proportion of welfare recipients who are Black propose only modest increases in benefits. In Appendix Table A3, we confirm that this correlation is strong and highly significant ($\beta = -0.401$, $p < 0.01$).

On the right of Figure 1, we find that the association between beliefs and policy support for UI benefits is significantly weaker than it is for TANF, as anticipated given the different histories of these programs ($\chi^2 = 6.43$, $p = 0.01$). In this case, respondents who believe relatively few Black people receive UI benefits tend to want to increase benefits by roughly a quarter, while respondents who think that almost all UI recipients are Black still want to increase the benefit by almost 20%. Correspondingly, Appendix Table A3 reports only small, insignificant correlations between these beliefs and support for UI ($\beta = -0.095$, ns).

C. Correcting misbeliefs

We argue that for information provision to be policy relevant, treatment effects must persist significantly beyond the intervention. Therefore, Stage 3 of the experiment examined the extent to which posterior beliefs of Black TANF or UI utilization in the control treatments are similar to participant prior beliefs while

⁶Importantly, there is also no significant difference in the prior beliefs held by our participants in the two treatment conditions of either experiment.




FIGURE 1. THE RELATIONSHIP BETWEEN BASELINE RACIAL MISBELIEFS AND PROGRAM SUPPORT AMONG UNTREATED PARTICIPANTS AS ESTIMATED IN THE SPECIFICATION REPORTED IN APPENDIX TABLE A3. SOLID LINES ILLUSTRATE REGRESSING CONTROL GROUP PROPOSED PERCENTAGE CHANGES TO PROGRAM BENEFITS ON MISBELIEFS ABOUT THE NUMBER OF RECIPIENTS WHO ARE BLACK. DOTTED LINES INDICATE ACCURATE BELIEFS. RESULTS FOR TANF (LEFT) AND FOR UI (RIGHT).

the posteriors of corrected participants are updated in the direction of the true statistics. To empirically assess the amount of persistence in uncorrected beliefs and attenuation in corrected misbeliefs, we estimate:

$$(Belief_i^{+1} - True_i) = \gamma_0 + \gamma_1 T_i + \gamma_2 (Belief_i - True_i) + \gamma_3 T_i (Belief_i - True_i) + \phi' \mathbf{X}_i + \epsilon_i$$

Where the left side is the Stage 3 posterior misbelief for participant i in either the TANF or UI experiment, T_i is an indicator for the participants who receive the information about the correct number of Black people benefiting from a policy in Stage 2, the difference $(Belief_i - True_i)$ is the respondent's Stage 2 prior misbelief, and \mathbf{X}_i is a vector of controls. In this case, γ_2 measures the stability or persistence of misbeliefs in the control population and the coefficient γ_3 measures the “pass through” of information provision. Here, we posit that $\gamma_3 \leq 0$. That is, the larger the prior misbelief in Stage 2, the smaller the posterior misbelief will be in Stage 3 if corrected participants update their beliefs toward the new information.

Figure 2 illustrates these estimates and reveals that participants in both experiments respond like Bayesian updaters. For participants in the control conditions, misbeliefs separated by a month are positively correlated and these associations are large and highly significant ($\gamma_2 = 0.478$, $p < 0.01$ for TANF and $\gamma_2 = 0.413$, $p < 0.01$ for UI). Hence, the misbeliefs of control respondents persist, on average, for at least a month. As important, we also see in Figure 2 that the information

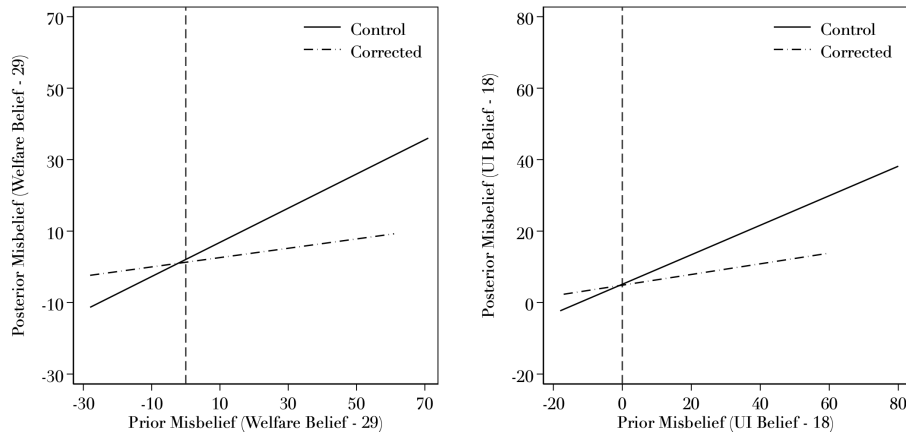


FIGURE 2. THE PERSISTENCE OF MISBELIEF CORRECTION ONE MONTH FOLLOWING THE INFORMATION INTERVENTION. THIS FIGURE SHOWS THE REGRESSION COEFFICIENTS γ_2 (SOLID LINE) AND $\gamma_2 + \gamma_3$ (DASH-DOT LINE) FOR TANF (LEFT) AND UI (RIGHT).

we provided about the true usage of TANF and UI, made a lasting impact on the beliefs of our participants in the treatment condition of both experiments. Here, we see that γ_3 is negative, as hypothesized, and the effects are sizeable. From Figure 2, we see that the linear estimates intersect very near a posterior misbelief of 0 (i.e., participants with correct pre-intervention beliefs retained those beliefs) and the relatively flat slope of the estimates for corrected respondents indicates that, regardless of the magnitude of their initial misbeliefs, after a month treated participants respond with a posterior belief much closer to the true statistic.

The magnitude and significance of the effects of providing this information on posterior misbeliefs are confirmed in Table A4, where we find $\gamma_3 = -0.348$ ($p < 0.01$) for the TANF experiment and $\gamma_3 = -0.263$ ($p < 0.01$) for the UI experiment. Lastly, Table A4 also indicates that adding controls changes these estimates very little. Summing, we find that the misbeliefs of respondents who did not receive any new information (in the control treatments) reported posteriors a month later that were similar and highly correlated with their priors. By contrast, respondents who did receive information to correct their prior misbeliefs largely believed it, report considerably smaller posterior misbeliefs, and this treatment effect lasts at least a month.

III. Discussion and Conclusion

In this paper, we share strong evidence that variation in beliefs about the racial composition of beneficiaries predicts support for welfare (TANF) but not unemployment insurance (UI) policies. We believe we are the first to compare

the programs in this setting—much of the literature focuses on welfare alone—and attribute the difference to the stigmatized racial history of welfare and the problematic notion of deservingness.⁷ We also show that beliefs about policies are persistent, at least on a month-to-month basis, and in subsequent work (Carpenter et al., 2024a), explore some important heterogeneities in beliefs, the elasticity of beliefs with respect to new information, and the implications of these beliefs for policy preferences.

Carpenter et al. (2024a) also extends our analysis in two important ways. First, when we induce persistent changes in beliefs, another striking difference between welfare and unemployment is observed—that is, beliefs about the composition of TANF recipients emerge as a causal determinant of support for welfare, but beliefs about UI do not cause support for unemployment insurance. This is in sharp contrast with Akesson et al. (2022), in which support for welfare is unaffected, and we speculate that the contrast reflects differences in methods of belief elicitation. In our case, the distribution of misbeliefs is roughly symmetric around zero. Second, we also investigate in Carpenter et al. (2024a) whether the induced changes in policy support are performative, and show that beliefs also cause donations to policy-related causes.

REFERENCES

- Akesson, Jesper, Robert W Hahn, Robert D Metcalfe, and Itzhak Rasooy.** 2022. “Race and Redistribution in the United States: an Experimental Analysis.” National Bureau of Economic Research.
- Alesina, Alberto, Matteo F Ferroni, and Stefanie Stantcheva.** 2021. “Perceptions of Racial Gaps, Their Causes, and Ways to Reduce Them.” National Bureau of Economic Research.
- Callaghan, Bennett, Leilah Harouni, Cydney H Dupree, Michael W Kraus, and Jennifer A Richeson.** 2021. “Testing the Efficacy of Three Informational Interventions for Reducing Misperceptions of the Black–White Wealth Gap.” *Proceedings of the National Academy of Sciences*, 118(38): e2108875118.
- Carpenter, Jeffrey, Jakina Debnam Guzman, Peter Hans Matthews, and Erin L. Wolcott.** 2024a. “Does Racial Animus Determine Redistributive Policy in the United States?” *Working Paper*.
- Carpenter, Jeffrey, Jakina Debnam Guzman, Peter Hans Matthews, and Erin Wolcott.** 2024b. “Racial Animus and Support for Labor Market Policies.” *AEA RCT Registry*.

⁷Ronald Reagan’s trope of the “welfare queen”, the undeserving (Black) mother with many aliases abusing government assistance, is perhaps the most famous example.

- Fong, Christina M.** 2007. "Evidence from an Experiment on Charity to Welfare Recipients: Reciprocity, Altruism and the Empathic Responsiveness Hypothesis." *The Economic Journal*, 117(522): 1008–1024.
- Gilens, Martin.** 1995. "Racial Attitudes and Opposition to Welfare." *The Journal of Politics*, 57(4): 994–1014.
- Haaland, Ingar, and Christopher Roth.** 2023. "Beliefs about Racial Discrimination and Support for Pro-Black Policies." *Review of Economics and Statistics*, 105(1): 40–53.
- Haaland, Ingar, Christopher Roth, and Johannes Wohlfart.** 2023. "Designing Information Provision Experiments." *Journal of Economic Literature*, 61(1): 3–40.
- 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.
- McGhee, Heather.** 2022. *The Sum of Us: What Racism Costs Everyone and How We Can Prosper Together*. One World.
- Quadagno, Jill S.** 1994. *The Color of Welfare: How Racism Undermined the War on Poverty*. Oxford University Press.
- Robbett, Andrea, and Peter Hans Matthews.** 2018. "Partisan Bias and Expressive Voting." *Journal of Public Economics*, 157: 107–120.
- Wolcott, Erin L.** 2024. "Did Racially Motivated Labor Policy Reverse Equality Gains for Everyone?" Federal Reserve Bank of Minneapolis 90.

Can Incorrect Beliefs about the Racial Composition of Welfare and Unemployment Insurance Beneficiaries be Changed?

Jeffrey Carpenter, Jakina Debnam Guzman, Peter Hans Matthews,
and Erin L. Wolcott

Online Appendix

TABLE A1—TREATMENT BALANCE ON OBSERVABLES.

	<i>TANF control</i>		<i>TANF treatment</i>		<i>UI control</i>		<i>UI treatment</i>	
	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
Age	42.570	15.004	43.632	15.331	42.667	14.570	42.99	15.134
Female	0.523	0.500	0.502	0.500	0.520	0.500	0.505	0.500
White	0.700	0.459	0.752	0.432	0.715	0.452	0.710	0.454
Black	0.124	0.330	0.105	0.306	0.137	0.345	0.131	0.337
College	0.433	0.496	0.424	0.494	0.411	0.492	0.420	0.494
Masters or more	0.159	0.366	0.167	0.373	0.157	0.364	0.156	0.363
Income over \$75k	0.408	0.492	0.436	0.496	0.442	0.497	0.394	0.489
Observations	709		707		706		703	

Notes: The F statistic from: (1) regressing a TANF indicator on the observables is 1.16 ($p = 0.32$),
(2) regressing an UI indicator on the observables is 0.76 ($p = 0.62$).

TABLE A2—DETERMINANTS OF PARTICIPANT BELIEFS AND MISBELIEFS.

	(1)	(2)	(3)	(4)
	Welfare Belief	Absolute Value of Welfare Misbelief	Unemployment Belief	Absolute Value of Unemployment Misbelief
Belief Confidence	-0.215 (0.199)	0.232 (0.122)	-0.261 (0.171)	-0.014 (0.127)
Age	-0.122 (0.037)	0.045 (0.022)	-0.147 (0.033)	-0.052 (0.024)
Female	0.014 (1.073)	1.361 (0.676)	1.030 (0.968)	2.366 (0.705)
College Degree	1.328 (1.094)	-0.709 (0.692)	-1.118 (0.972)	-1.339 (0.714)
Income over \$75k	1.022 (1.093)	0.857 (0.681)	-0.787 (0.975)	-0.638 (0.717)
South	1.933 (1.209)	0.845 (0.754)	0.752 (1.094)	0.463 (0.808)
White	2.110 (1.262)	0.655 (0.784)	-1.939 (1.168)	-0.474 (0.869)
IAT (biased against Blacks)	1.777 (1.098)	1.428 (0.685)	0.527 (0.989)	0.377 (0.723)
Explicit (preference for Whites)	1.731 (1.224)	0.099 (0.762)	2.845 (1.116)	1.409 (0.816)
Constant	31.628 (2.061)	10.365 (1.260)	30.489 (1.934)	14.665 (1.444)
Dependent variable sample mean	30.08	15.66	22.95	13.41
Observations	1383	1383	1368	1368

Dependent variables “Welfare Beliefs” and “Unemployment Beliefs” are respondents’ prior beliefs (from Stage 2) about the number of Black people, out of 100, using each program. Dependent variable “Absolute Value of Welfare Misbelief” is $|Welfare\ Belief - 29|$ and dependent variable “Absolute Value of Unemployment Misbelief” is $|UI\ Belief - 18|$. OLS with robust standard errors reported.

TABLE A3—PRIOR BELIEFS AND POLICY SUPPORT.

	(1)	(2)	(3)	(4)
	Welfare	Welfare	Unemployment	Unemployment
Prior Misbelief	-0.401 (0.085)	-0.393 (0.086)		
Prior Misbelief			-0.095 (0.086)	-0.121 (0.088)
Constant	35.605 (1.584)	38.912 (5.373)	22.473 (1.337)	32.307 (4.760)
Controls	No	Yes	No	Yes
Dependent variable sample mean	34.00	34.00	21.77	21.77
Observations	709	707	710	706

The dependent variable is a continuous measure ranging from -100 to 100 which indicates the percentage by which the respondent would like to change TANF or UI support. Controls include age, sex, education, income and geographic region. OLS with robust standard errors reported.

TABLE A4—MISBELIEF CORRECTION AND PERSISTENCE.

	(1)	(2)	(3)	(4)
	Welfare	Welfare	Unemployment	Unemployment
Treated	-0.783 (0.990)	-0.776 (0.989)		
Prior Misbelief	0.478 (0.040)	0.479 (0.040)		
Treated × Prior Misbelief	-0.348 (0.058)	-0.351 (0.058)		
Treated			-0.256 (0.932)	-0.308 (0.920)
Prior Misbelief			0.413 (0.047)	0.406 (0.047)
Treated × Prior Misbelief			-0.263 (0.062)	-0.257 (0.062)
Constant	2.070 (0.697)	3.603 (1.773)	5.131 (0.662)	6.914 (1.735)
Controls	No	Yes	No	Yes
Observations	1161	1159	1163	1155

Dependent variable is posterior misbelief (from Stage 3). OLS with robust standard errors reported.

Controls include age, sex, education, income and geographic region.