

Popular Support for Equity in Algorithmic Decision-Making: Case Study in Online Welfare Ads

Keywords: algorithmic fairness, online advertising, preference elicitation, welfare programs, political party identity

Introduction: Algorithmically optimizing the provision of limited resources is commonplace across domains from healthcare to lending. Optimization can lead to efficient resource allocation, but, if deployed without additional scrutiny, can also exacerbate inequality. Little is known about popular preferences regarding acceptable efficiency-equity trade-offs, making it difficult to design algorithms that are responsive to community needs and desires. Here we examine this trade-off and concomitant preferences in the context of GetCalFresh, an online service that streamlines the application process for California’s Supplementary Nutrition Assistance Program (SNAP, formerly known as food stamps), which helps low-income people buy nutritious food. GetCalFresh runs online advertisements to raise awareness of their multilingual SNAP application service. We first show via experiments that when online ads are optimized to garner the most enrollments per dollar, a disproportionately small number of Spanish speakers enroll due to relatively higher costs of non-English language advertising. We next embed these results in a survey (N = 1,532) of a diverse set of Americans, finding broad popular support for valuing equity in addition to efficiency: respondents generally preferred reducing total enrollments to facilitate increased enrollment of Spanish speakers. These results buttress recent calls to reevaluate the efficiency-centric paradigm popular in algorithmic resource allocation.

GetCalFresh Ads: Roughly 12% of Americans annually participate in SNAP [2], though many more low-income individuals are eligible. One barrier to participation is that the sign-up process can be confusing and time-intensive. Spanish speakers make up 23% of individuals under the poverty line in San Diego; there is often a language barrier among this population when applying for SNAP benefits. The online service GetCalFresh aims to address this obstacle by simplifying SNAP enrollment and recruiting individuals through online ads.

Focusing on Google Search Ads in San Diego County, we experimentally observe significantly higher cost-per-conversion for Spanish language SNAP ads relative to English language SNAP ads. Depending on the algorithmic bidding mechanism used, the average daily cost per conversions for **a Spanish-targeting ad campaign is anywhere from 1.4 to 3.8 times more costly** than for an English-targeting ad campaign. We postulate that the reason for these language-based disparities could stem from the demand side: fewer Spanish speakers are Googling for these keywords relative to English-speaking counterparts.

Our online ad experiments reveal stark disparities in the cost of recruiting English and Spanish speakers into SNAP. Those cost disparities in turn create an inherent tension between efficiency and equity: reaching more Spanish speakers means lowering the overall reach of the advertising campaign, as Spanish speakers are, on average, more expensive to recruit (Figure 1). Another research question arises: how can GetCalFresh allocate their advertising budget to ensure its SNAP enrollees include a “fair” share of Spanish speakers?

Preference Elicitation: To understand preferences for SNAP ad budget allocation between English and Spanish speakers, we ran a Qualtrics survey distributed via Prolific to a gender-balanced sample of U.S.-based partisans (Republicans or Democrats). Respondents were shown a series of pairwise comparisons of feasible ad budget allocation outcomes along the Pareto

frontier from our online experiment described above. For each pairwise comparison, respondents were asked to choose their preferred allocation (e.g., choosing between either Option A or Option B in Figure 1).

We then performed conjoint analysis on surveyed preferences to measure the extent to which individuals are willing to exchange some degree of efficiency in order to reach more Spanish speakers in the SNAP online advertising paradigm. We define “efficiency” here as maximizing the number of conversions per dollar, which corresponds to selecting the lowest-possible allocation of Spanish speakers (since Spanish speakers are, on average, more expensive to target than English speakers). Overall, we find **low preference for the most “efficient” option across respondent demographic groups**—across age, race, gender, education, income, religion, and party identity—with the largest point estimate preference difference (of 20 percentage points) being between Republicans and Democrats. Figure 2 depicts poststratified logistic regression estimates of the share of a demographic group’s preference for the most efficient allocation.¹ To understand preferences for deviating from the most efficient allocation, we study the win rates of each budget allocation option based on our survey’s pairwise comparisons. In Figure 3, we see that the win rates for the most-efficient allocations are nearly always far lower than (if not comparable to) the win rates near “demographic parity” of 23% Spanish speaker conversions—even among surveyed Republicans.

Algorithmic Fairness: Our findings on preference for equity stand in contrast to general population preferences for affirmative action in the workplace and in college admissions. Within our survey, the share of respondents preferring to prioritize Spanish speakers over a purely “efficient” advertising strategy is far higher than those supporting affirmative action policies; this points towards broader support for equity in machine learning-based algorithmic applications, which has significant implications for future technology policy. Our results have immediate implications for the equitable design of algorithms. Deployed optimization algorithms—like those used in online advertising—largely focus on efficiency, but our results reveal that when presented with detailed equity-efficiency trade-offs, the general public would incorporate equity considerations into allocation decisions. As a result of this research, **GetCalFresh has adjusted their Google Ads budget allocations** to recruit more Spanish speakers, roughly in line with the demographic parity benchmark (e.g., 23% Spanish speakers in San Diego County). Furthermore, whereas much of the algorithmic fairness literature has focused on context-agnostic, axiomatic approaches, **our results illustrate the value of framing questions of equity in terms of concrete trade-offs that decision-makers, affected communities, and the general public can consider** when making difficult choices.

References

- [1] Trevor Tompson and Jennifer Benz. Ap votecast 2018. *Ann Arbor, MI: Inter-university Consortium for Political and Social Research*, 2019.
- [2] USDA Food and Nutrition Service. Snap data tables, August 2020.
- [3] Wei Wang, David Rothschild, Sharad Goel, and Andrew Gelman. Forecasting elections with non-representative polls. *International Journal of Forecasting*, 31(3):980–991, July 2015.

¹We first model the preferred options conditional on respondent demographic data on gender, age, party identity, race, religion, education, and income; we then estimate the preferred option on each of 3,840 cells (for each combination of demographic subgroups). Finally, we take a weighted average of cell-level estimates, with weights equal to the share of the U.S. population belonging to each cell based on 2018 AP VoteCast data [1, 3].




Figure 1: The three options displayed each represents a feasible allocation of the same \$400 daily advertising budget; Option A shows the most efficient allocation (> 0% Spanish conversions since individuals targeted with English ads may still be Spanish speakers) and Option B shows an allocation aligning with “demographic parity” for San Diego.

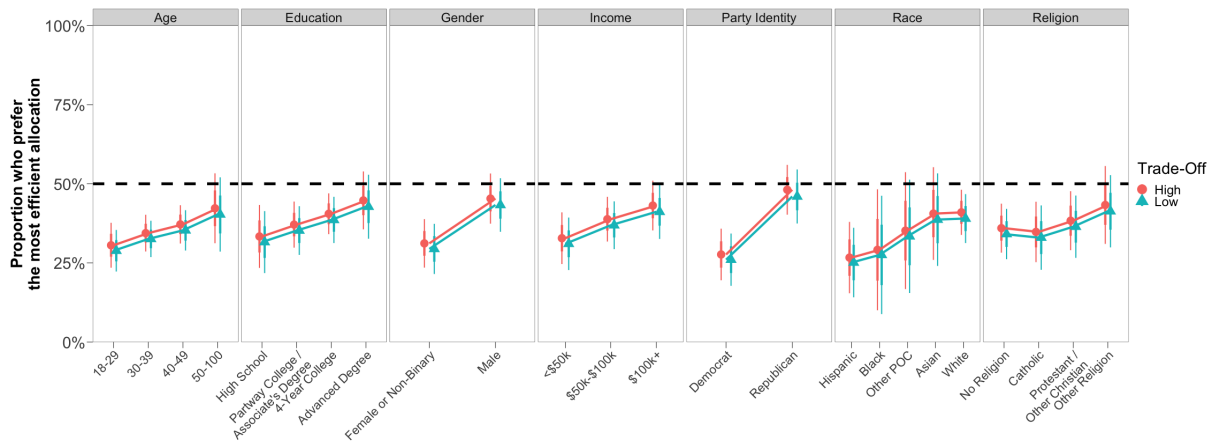


Figure 2: A minority of respondents prefer the most efficient allocation across all demographic groups; results are shown for Pareto frontiers generated using two different bidding algorithms yielding either high or low trade-offs between English and Spanish costs per conversion.




Figure 3: The “most efficient” allocation is not the most popular allocation for US Democrats or Republicans; in most observed settings, the allocation closest to “demographic parity” (23% Spanish speaker conversions) was most popular.