Generosity Across the Income and Wealth Distributions

Jonathan Meer*
Texas A&M University
and NBER

Benjamin A. Priday
Texas A&M University

April 2020

Abstract

Despite widespread interest, there is little systematic evidence on the relationship between income, wealth, and charitable giving. We use the Panel Study of Income Dynamics to provide descriptive statistics on this relationship. We find that, irrespective of specification, donative behavior increases with greater resources.

*Meer: jmeer@tamu.edu. Priday: priday@tamu.edu. We are grateful for comments from Dan Hungerman, Judd Kessler, Benjamin M. Marx, Harvey Rosen, and Mark Wilhelm. Participants at the 2019 Science of Philanthropy Initiative Conference provided valuable feedback as well.

1. Introduction

The relationship between income, wealth, and charitable giving is of significant policy and popular interest, yet there is surprisingly little systematic evidence on the topic. While the press tends to highlight findings that suggest that the well-off are stingy (see, among many others, Grewal, 2012, and Young, 2017), the claims are questionable due to misuse of data, incomplete controls, inappropriate empirical specifications, and a lack of accounting for the influence of outliers.[1]

Using the Panel Study of Income Dynamics, we provide the most rigorous documentation of these relationships to date. We estimate the relationship between pre-tax income, wealth, and charitable giving over three measures of generosity: the likelihood of giving, the amount given, and the percent of income given. We report means, estimates using ordinary least squares with a rich set of household demographic controls, as well as estimates with household fixed effects. Our statistical models take care to mitigate the potentially misleading impact of outliers by winsorizing households' percent of income given. This allows us to simultaneously retain the full sample while limiting the degree to which outliers can skew estimation results. While the PSID undersamples very-high-income households, it provides the most complete view of charitable giving across the income distribution, and includes

[1] As an example, Stern (2013) documents the alleged miserliness of the rich, as measured by the percent of income given, using analysis from the Chronicle of Philanthropy discussed in Gipple and Gose (2012). However, that study was based on data from the Internal Revenue Service, for households that itemize their giving. Tax-filing units with incomes below $50,000 were not included, comprising nearly two-thirds of households. Thus, referring to giving by the "poor" is misleading. More importantly, though, itemizing one's deductions is an endogenous decision. At the lower levels of the income range in the study, only about half of tax returns were itemized, while at higher levels, nearly all returns were. Almost by definition, those who itemize despite lower incomes have higher levels of charitable giving; otherwise, they would not itemize. Further, changes to the tax code under the Tax Cut and Jobs Acts of 2017 will significantly reduce the number of households that itemize, making tax data even more problematic for examining charitable giving (Meer and Priday, 2019).

data on wealth, religiosity, and a number of other important correlates of giving that are rarely present in other data.[2]

We find that, irrespective of the specification, higher-income and higher-wealth individuals are substantially more likely to make donations to charity, and give significantly more. Once outliers are properly accounted for, the relationship between the percent of income given and income or wealth is generally flat. We also analyze data from the Internal Revenue Service Statistics of Income, and find that the very-highest-income individuals tend to give the most as a percent of their income.

One common finding in the literature is the "U-shaped" giving curve: the average proportion of income donated is largest at the lowest and highest parts of the income distribution, while middle-income groups give the smallest proportion of their income. For example, Clotfelter and Steuerle (1981) use IRS tax return data and find a U-shaped curve, which is later corroborated with a panel of tax filers by Auten, Clotfelter, Schmalbeck (2000), though only with mean, not median, giving.[3] However, there is no consensus on this point. Wiepking (2007) documents the relationship between income and giving in the Netherlands, finding a flat relationship for the likelihood of giving and a negative one on the percent of income given. Schervish and Havens (1995a, 1995b) estimate that the giving curve is essentially flat when considering the full sample, not just those who donated a positive amount. James and Sharpe (2007) argue that selection bias drives the Schervish and Havens results. They find a U-shaped curve but note that it is not a valid description of the typical household within each income level: older "low-income" donors with higher wealth are responsible

[2] Wealth, in particular, is an important determinant of giving (James and Sharpe, 2007), particularly at the end-of-life (Bakija, Gale, and Slemrod, 2003; Meer and Rosen, 2013). Households may give not just out of income, but from existing wealth. As we show below, a good deal of the ostensibly higher giving levels by low-income households are in fact driven by those with high levels of wealth.

[3] Few studies use panel data, and tax return data must be limited to those who itemize in every year. Rooney, Wilhelm, Wang, and Han (2019) show that the pattern of charitable giving across individuals when using a single year (or repeated cross-section) can be misleading.

for the higher mean giving at the bottom of the income distribution. Importantly, other than James and Sharpe (2007), the effect of outliers has largely been ignored in this literature which, as we show below, leads to incorrect inference.[4]

Experiments, both in the lab and the field, have tackled this question as well. Andreoni, Nikiforakis, and Stoop (2017) summarize the literature on the relationship between prosocial behavior and income and test it by misdelivering envelopes in higher and lower resource neighborhoods. They vary the amount of money in the envelope and the ease with which the false recipient can retain it, and find that the rich are more than twice as likely to return the envelope as the poor. However, they argue that these results reflect different financial pressure and marginal utility of income, and that social preferences are in fact similar. Lab experiments are limited in the external validity of their results, as they rely either on the existing (and endogenous) level of resources of the subjects (Blanco and Dalton, 2019), or differences in "income" of just a few dollars created during the experiment (Duquette and Hargaden, 2019).

But there is a broader set of questions underlying this discussion: what is generosity, and how should it be measured?[5] Even taking a limited view – primarily for data-availability reasons – that generosity is measured by donations of money, the metric remains unclear. Is it the likelihood of making any donation at all? The amount given? The amount given as a share of one's income? Of one's total resources, measured as wealth or permanent income? Over what period of time?

Ultimately, the question of what it means to be generous is philosophical and deeply normative. Whether one believes that the consequences of behavior (that is, observed giving)

[4] Duffy *et al.* (2015) use the 2001-2009 waves of the PSID to look at how income and wealth affect giving, with special attention given to directly comparing the PSID to James and Sharpe (2007).

[5] Gee and Meer (2019) discuss a number of related issues surrounding the notion of an "altruism budget." Lilley and Slonim (2014) and Brown *et al.* (2019) analyze the substitution between gifts of time and money.

determines this state, or whether underlying intentions are more important (as Andreoni, Nikiforakis, and Stoop, 2017, do) is a personal matter. We strive to provide a set of facts about the empirical relationship between income, wealth, and giving, and leave it to the reader to draw his or her own conclusions.

Different empirical approaches provide insight into very different questions. A simple comparison of means (or medians) across the income and wealth distribution shows the most straightforward relationship between resources and giving, but compares people who face very different circumstances and are at different stages of life. That is, much of the observed relationship may be due to other factors that are correlated with both income (or wealth) and giving, so this approach shows only whether the sorts of people who have more money are the sorts of people who donate more. Ordinary least squares estimates examine the relationship conditioned on the effect of these factors, purporting to measure the effect of resources on giving, holding all else equal. But *unobserved* differences between people might still be driving the relationship. Using the panel nature of the PSID and including household fixed effects accounts for time-invariant characteristics, like altruism, and allows us to estimate how changes in income or wealth relate to changes in donative behavior, taking into account both observed and unobserved individual characteristics. This approach comes closest to answering how a household's charitable giving changes when it has more resources. However, as we detail below, even these results should not be viewed as completely causal.

2. Data and Specification

2.1 Data

The Panel Study of Income Dynamics (PSID) is a biennial household survey that collects information from the previous year on wealth, income, and a rich set of individual and household characteristics. The Center on Philanthropy Panel Study (COPPS), a PSID

module begun in 2001, includes data on charitable giving to religious entities and ten categories of secular charities. We aggregate household donations for our primary analysis and explore possible differences in giving to religious and secular causes later.

The PSID is especially advantageous for analyzing charitable activity. Wilhelm (2006) shows that the data are superior to other surveys in coverage of charitable giving and compare well to tax return data below relatively high levels of income. Tax return data cover high earners – who constitute a large portion of charitable giving – and have higher precision; however, filers only appear in the data if they itemize contributions. This effectively results in a sample selected for higher income levels, precluding estimation across the broader income distribution. Recent changes in the tax code have dramatically reduced the number of households that itemize, limiting the value of these data moving forward (Meer and Priday, 2019). Tax return data also do not disaggregate giving by charity type. Further, data on wealth are not available on tax returns, although wealth is an important determinant of donative behavior (Bakija and Heim, 2011; James and Sharpe, 2007). Below, we show that wealth is important to consider when estimating giving behavior across all levels of the income distribution, not just high earners.

We use nine waves of the PSID spanning 2001–2017, representing every other calendar year from 2000 to 2016. The raw sample consists of 76,834 observations across 16,148 households. We remove observations with negative income (50) as well as the 2017 Immigrant Refresher (452) and the Survey of Economic Opportunity low-income oversample (23,199).[6] This leaves us with a final sample of 53,133 observations over 10,665 households.

[6] The immigrant refresher added new households to the panel beginning in 2017 to ensure the data remain reflective of the current demographic composition of the country. We omit this group as they appear in only one year of our data. The SEO subsample was included in the PSID to allow researchers to focus on low-income households. Including it in our analysis would oversample the bottom of the income distribution, skew bin specifications, and likely bias results. Given that the panel is nationally representative without the oversample, we opt to simplify analysis by removing these households instead of weighting the data.

Across all household-years, 60 percent of households report making a donation. Conditional on making a donation, the mean (median) gift is $2,749 ($1,000) in 2018 dollars. The mean (median) income of our sample is $91,646 ($67,620); the mean (median) wealth is $356,082 ($71,400).[7] Table 1 contains the full set of summary statistics.

We examine three outcomes: the probability of giving, the log amount donated, and the percent of income donated, focusing on the first two. Outliers pose a particular problem when comparing the percent of income given across income levels. James and Sharpe (2007) show that the U-shaped giving curve is almost entirely driven by a small number of outliers at the bottom of the income distribution. These are often elderly households with high levels of wealth but low income. Non-elderly households with transitorily low income, perhaps through business losses, may maintain giving behavior of other, higher-income years. Such "low-income" donors can distort means, as we show below. Controlling for wealth and age helps mitigate this problem but is not sufficient: outliers may still disproportionately skew means and mislead inference.

James and Sharpe (2007) address this by simply removing outliers with giving above a threshold percent of income. However, this introduces selection bias that is especially problematic when utilizing panel data; removing a household's observation in one year changes within-household estimation. Instead, we winsorize the percent of income given, a not-uncommon approach in other fields but relatively novel in charitable giving (Eckel, Herberich, and Meer, 2018). We cap the percent of household family income donated at 20 percent: any household-year observations donating more than that are assigned a value of 20 percent for percent of income given, while all other observations remain unchanged. We

[7] The PSID defines "family income" as the sum of taxable, transfer, and social security income for all family unit members. Wealth is the sum of a household's imputed business and farm equity, transaction accounts, value of debt other than main home mortgage or vehicle loans, home equity, other real estate equity, stock equity, vehicle equity, IRAs, and other assets.

also generate the winsorized amount given by restricting the maximum amount donated by a household-year to be 20 percent of income.[8] These approaches allow us to retain all observations while limiting the disproportionate impact of outliers on intensive margin outcomes. Table 2 compares winsorized outcomes to their non-winsorized counterparts across income quantiles. While about 6 percent of the observations in the bottom income ventile are affected by winsorizing at 20 percent of income, this is sufficient to reduce the mean percent of income given by a factor of 16.

2.2 Specification

We divide family income and wealth into bins for analysis. Income is broken down into evenly-sized ventiles. Because there are a number of observations with negative or zero wealth (7,346 and 1,650, respectively), we create separate bins for each of these and divide all positive wealth into twenty evenly-sized bins, resulting in twenty-two wealth bins. Table 3 lists the range of the income and wealth bins.

Discretizing income and wealth into bins reduces the effect of outliers in estimation without imposing a functional form on the relationship. However, using bins means that marginal effects measure movements across the income or wealth distribution, rather than the change from an additional dollar of income or wealth.

There are well-documented determinants of giving behavior including age, gender, education, and race (Bekkers and Wiepking, 2011). We include controls for these characteristics, as well as marital status, retirement and disability status, self-reported health, and religious affiliation; number of children, a housing price index and its quadratic; as well as state and year effects.[9]

[8] For robustness, we consider 10%, 20%, and 30% winsorization when analyzing the percent of income given. The log winsorized amount given uses a 20% winsor for all analysis.

[9] Time-invariant characteristics are excluded in fixed effects specifications. Housing price and its quadratic are included to take into account the macroeconomic environment (List and Peysakhovich, 2011; Meer, Miller, Wulfsberg, 2017).

Addressing corner solutions – that is, non-donors – is a classic problem in the charitable giving literature. The Tobit is frequently used to account for this problem. However, this model suffers from tractability problems in the presence of fixed effects, is likely not appropriate when zeroes arise from corner solutions rather than true data censoring, and constrains the marginal effects on the extensive and intensive margins to be proportional to each other. This last issue is particularly problematic when considering the impact of, say, income, which may have very different impacts – possibly of opposite signs – on the likelihood of making a donation and the amount given.

We follow a more recent trend in the literature and treat giving as a two-step process: a household first decides if they will give, then, if so, how much they will give. Using ordinary least squares (OLS) and fixed effects (FE), we estimate the impact of income and wealth on the extensive and intensive margins of giving, then combine the results to find marginal effects on the unconditional amount given.[10] Standard errors are clustered at the household level in all specifications.

It is important to note that OLS and FE regressions answer different questions. Marginal effects from OLS give the impact on giving of moving across the income and wealth distributions, holding all observable characteristics equal. However, unobservable factors – particularly altruism – are likely to vary across these distributions, making it difficult to view these results as causal. If, for example, the types of people who hold large amounts of wealth tend to be less generous, even taking into account observable characteristics, then the giving-wealth gradient will be negative; it would be incorrect, however, to conclude that higher levels of wealth *makes* people less generous. These specifications, therefore, are better

[10] See Huck and Rasul (2011) and Meer (2011) for more discussion on the use of this specification for estimates of charitable giving responses. Wilhelm (2008) compares the Tobit to other specifications for charitable giving estimation.

suited to shed light on whether people who have or earn more money tend to be more or less generous than those who are observationally similar, but who have less money.

Fixed effects estimates hold time-invariant characteristics, like innate altruism or permanent income, fixed; that is, they measure how within-household changes in income or wealth impact within-household changes in giving. These estimates are more suited to answer questions about whether people become more or less generous as they acquire more money, and the marginal propensity to donate out of income or wealth. With controls for wealth and (implicitly) permanent income, these should also be viewed as the propensity to donate out of *transitory* income.

However, these fixed effects estimates require a nontrivial assumption to be taken as causal, namely, that *changes* in unobserved factors are uncorrelated with financial resources and giving. For instance, if a manager becomes more willing to help others both in her work and personal life, we might observe that she receives a pay increase and that her charitable giving increases. We would mistakenly conclude that the former causes the latter, even though both were caused by the change in her attitude and behavior. As such, even fixed effects estimates cannot be viewed as being a fully causal accounting of the relationship between income, wealth, and giving. To estimate a causal relationship in a compelling way, one would need random (or as-good-as-random) changes in income and wealth that are completely unrelated to individuals' own choices and (unobservable) attributes. Only then could we be assured of observing a causal impact on giving. Needless to say, it is quite difficult to imagine such circumstances. Further, there is evidence that giving behavior varies by the type of income shock (Drouvelis, Isen, and Marx, 2019); varying individuals' labor income or financial assets at random seems even more far-fetched than observing windfall gains at a large scale.

2.3 Limitations

The PSID has very few observations on high-income households, which make up the majority of donations and likely behave differently than households in our sample; we cannot directly infer the effects on their generosity using our results.[11] Tax filing units with earnings over $2 million made $63 billion in deductible contributions in 2016, or about 30 percent of the total, despite making up about 0.1 percent of the population (IRS, 2018). But there are only two observations with income above that level in the PSID in that year.

Additionally, the PSID samples households every two years and therefore depends on a respondent's ability to recall information from the previous year. Respondents may lie or forget about their income, wealth, or charitable donations, though Wilhelm (2006) argues the PSID is reliable and superior to other survey data in this respect. The rising popularity of donor-advised funds means people may shift charitable giving across years (Andreoni, 2018), which might be missed in our biennial data. Our data also covers over a decade of changes in the tax treatment of charitable donations; while we do not directly consider the effect of tax laws here, virtually all studies find that they do affect giving (see Meer and Priday, 2019, for recent evidence on that question).

The percent of income given is an outcome of particular interest in existing literature and popular media. But regression analysis of this outcome has a serious problem: the dependent and independent variables both contain family income. It is somewhat strange to discuss the impact of income on the percent of income donated while assuming that the denominator of donations over income is held constant.

[11] A handful of studies have examined giving behavior by high-income prospective donors using field experiments, though few consistent conclusions have emerged. See Alston *et al.* (2018); Kessler, Milkman, and Zhang (2019); and Levin, Levitt, and List (2016).

3. Results

3.1 Means and Outliers

We begin by examining mean giving behavior by income (or wealth) bin. This approach does answer the simple question of whether giving by higher-resource households differs from lower-resource households, but without accounting for any characteristics that might be different between those groups that drives philanthropy.

Figure 1 shows the probability of making a donation in a given year across the income distribution. The probability rises monotonically, from 22.3 percent in the lowest bin to 91.0 in the highest. The difference between the likelihood of making any donation between the bottom and 10th bins (37.0 percentage points) is nearly as large as the difference between the 10th bin and the highest one (31.6 percentage points).

The relationship for wealth in Figure 2 is slightly more complicated. Recall that the lowest bin is for those with negative wealth, while bin 2 comprises those with zero wealth. The negative-wealth bin has an average likelihood of giving of 45.3 percent, as compared to 13.4 percent for those with zero wealth. It is not until about the 25th percentile of the positive wealth distribution (bin 7) that the likelihood of giving is as high as for those in the bottom bin. This likely reflects the different demographics of those in the bottom bin, who are the youngest group on average, but are as likely to have a college degree or more as those around the 60th percentile of the wealth distribution. Among the ventiles of positive wealth, the likelihood of giving increases steadily from 26.8 percent to 90.3 percent.

Figures 3 and 4 show the mean amount given by income and wealth bin, both unconditionally and winsorizing at 20 percent of family income. Unsurprisingly, giving increases with higher levels of resources (excepting those with negative wealth). Winsorizing does not change most bins' mean giving very much: the top income bin is reduced to $6,481 from

$6,553. The lowest-income bins are affected most in relative terms; the bottom bin's winsorized mean is $108 as compared to $248 unconditionally, while the second bin's mean is reduced to $215 from $271.

Figure 5, which shows the mean percent of income given by bin, illustrates the sensitivity of results to the presence of outliers. The unconditional mean percent of income by those in the bottom bin is 33.6 percent, including zeroes for the nearly four-fifths of observations in the bin who do not give at all. For the other bins, the mean varies between 1.7 and 2.1 percent. Figure 6 shows percent of giving winsorized at 10, 20, and 30 percent of family income. It is clear that a small number of extreme observations drive the mean for the lowest income bin; indeed, the mean percent given conditional on making a donation is about 146 percent, and 16 observations have giving in excess of 500 percent. Giving as a percent of income tends to be fairly flat across the income distribution when the influence of these outliers is reduced. Table 2 shows that fewer than 6 percent of observations in that bottom bin are affected by winsorizing, but the difference in outcome and interpretation is vast.

Figure 7 illustrates the mean percent of income given by wealth bin. Once again, the influence of extreme outliers is visible: mean giving at the 25^{th} percentile of the wealth distribution is 18 percent of income, far higher than any other bin, and driven primarily by a single observation (from the lowest income bin). Figure 8 shows the winsorized percent of income given by wealth bin. Excluding the negative-wealth bottom bin, the relationship is strongly positive, rising monotonically from 0.3 percent of income to 3.5 percent of income in the top wealth bin.

As discussed in Section 2.3, one of the primary limitations of the PSID is the lack of data on very high-income households, who make up the bulk of the dollars given. Table 4 shows the mean percent of giving for those groups, based on the Internal Revenue Service

Statistics of Income in 2016, which report itemized giving. For simplicity (and as a conservative approach), we assume non-itemizers do not donate: in reality, non-itemizers are likely to donate some positive amount, which makes our estimates a lower bound. While the results are not directly comparable to those from the PSID, the relationship between income and the percent of income given remains strongly positive. We do include this calculation for the $200,000-$500,000 bin of adjusted gross income, which corresponds roughly to most observations in our top income bin; mean giving for that group is 1.97 percent, not dissimilar to the 1.72 percent we find in the PSID. The likelihood of giving for these groups (as defined by itemized giving) is between 86 and 94 percent. Higher income bins have higher averages, increasing from 2.3 percent for $500,000 to $1 million to 4.3 percent for $5 to $10 million, and 8.6 percent for those with adjusted gross incomes above $10 million.

Taken together, these results show the importance of accounting for outliers in this type of analysis, as well as the need to account for both income and wealth when considering the giving gradient. As noted above, however, these findings document only a strong positive correlation between resources and giving behavior. In the next section, we conduct regression analysis to examine the degree to which income and wealth affect giving, even after taking into account observable individual characteristics that may be correlated with giving.

3.2 Ordinary Least Squares Estimates

We examine the degree to which observable characteristics explain the income- and wealth-giving gradients beginning with Figure 9. Each graph includes three sets of coefficients. The first reports the marginal effect of being in an income (wealth) bin with no controls other than indicators for year, relative to the lowest income (wealth) bin. The second adds our full set of covariates: age and its quadratic, indicators for level of education, marital status, state of residence, religious denomination, retired and disabled status, indi-

cators for self-reported health status, race, gender, and a housing price index and its quadratic. The third set of results adds wealth (income) to the controls, thus examining how giving behavior varies across the income distribution holding both observable characteristics and wealth (income) constant.

If, for example, the strong positive relationship between income and giving discussed in Section 3.1 is merely a reflection of other characteristics that are correlated with income and giving (such as education or religious affiliation), then the second and third set of results will show a flat gradient. That is, the independent effect of income on giving, conditional on those observables, will be zero.

Figure 9 through 12 indicate that the relationship remains strongly positive, though these controls do account for a good proportion of the impact. In Figure 9, we show results for the likelihood of giving. Only including year effects, the difference between the top and bottom income ventiles is 68.2 percentage points, as expected from the differences in means. Adding observable characteristics maintains the monotonic relationship, but reduces this difference to 38.7 percentage points. That is, the likelihood of giving is much higher for individuals who are observably similar but have higher incomes, on average. Adding a set of indicators for wealth bin further flattens the relationship, though those in the top bin are still more than 27 percentage points more likely to make a donation than those in the bottom. The differences between the very top income bins are particularly reduced, with the propensity to make a donation roughly the same after the 80^{th} percentile of income.

Similarly, Figure 10 shows the relationship between wealth and the likelihood of giving. Interestingly, the higher likelihood of giving by those with negative wealth remains large and significant relative to those with low levels of wealth when including controls (though the difference is reduced by a large margin), suggesting that there is more to this correlation than simply age, education, and other observable characteristics. The top wealth bin is about

15 percentage points more likely to give than those with negative wealth (and 29 percentage points more likely to give than those with zero wealth) when including the full set of controls. Again, the relationship tends to flatten out past the 70th percentile of the positive wealth distribution.

In Figures 11 and 12, we turn to effects on the unconditional amount of given, winsorized at 20 percent of family income. As described in Section 2.2, we estimate the extensive and intensive margins separately and then combine them to obtain unconditional estimates. Once again, the income-giving gradient is strongly positive and, in this case, only somewhat attenuated by the inclusion of observable characteristics, including wealth. Those in the top income bin give about 16 times more, on average, than those in the bottom income bin; of course, the mean income in the top bin is 68 times that of the bottom bin. In Figure 12, the wealth-giving gradient is similarly positive, with those in top wealth bin giving about 380 percent more than those with negative wealth, and about 1300 percent more than those with zero wealth.

As discussed above, it is unclear how best to interpret the effects of income on the percent of income given. Nevertheless, for comparison with previous findings, we show these results by income and wealth bin, respectively, in Figures A.1 and A.2, winsorizing at 10, 20, and 30 percent of family income and including wealth and demographic controls. The results do not change meaningfully across these levels of winsorizing, suggesting that it is a small number of extreme outliers that drive the conclusions often seen in the popular press. The percent of income given increases with income, plateauing at about median income, before reducing somewhat for the very top income bin. The results for wealth are more clear: the percent of income given rises monotonically with wealth (excluding those with negative wealth), peaking at about two percentage points higher for the top wealth bin relative to those with negative wealth.

The findings are stark: generosity – as reflected by donations – increases with greater income and wealth. But these results could still reflect unobserved heterogeneity, necessitating stronger controls. We therefore now turn to estimates with individual fixed effects.

3.3 Fixed Effects Estimates

As discussed above, fixed effects estimates account for unobservable, time-invariant attributes, including innate altruism. While these are not necessarily consistent estimates of the underlying marginal propensity to donate out of additional income or wealth, they are closer to that interpretation than the OLS estimates. Suppose, for example, unobserved altruism was correlated with both income and giving but not with the observable characteristics listed in Section 3.2. Then those OLS estimates would still reflect spurious correlation, while estimates including fixed effects would show no relationship between resources and giving.

In Figures 13 and 14, we examine the effects of income and wealth on the likelihood of giving. The gradient remains very positive, as with OLS, though the differences between bins are unsurprisingly much smaller. For example, including time-varying covariates and wealth, an individual in the top income bin is 13 percentage points more likely to make a donation than that individual would be had he or she been in the bottom income bin. Note further that the difference between the three specifications is much smaller than with OLS, since the individual fixed effects account for so much of the variation picked up by the covariates in Figure 9. The relationship is similarly positive for wealth. In Figure 14, with those in the top wealth bin are 6.6 percentage points more likely to give than if they had been in the negative wealth bin, and 11.9 percentage points more likely than if they had zero wealth.

Figures 15 and 16 examine the effect on the unconditional amount given. Again, the gradient remains strongly positive, with those in the top income bin giving about 420 percent

more than they would have had they been in the bottom income bin. And the same holds for higher levels of wealth: those in the top wealth bin give, on average, 100 percent more than they would have had they had negative wealth.[12]

The fixed effects estimates suggest that unobserved heterogeneity is not entirely driving the positive relationship between income (wealth) and donative behavior. By these measures, individuals appear to become more generous as their resources increase, though we caution against making too strong of a causal statement on this front.

3.4 Religious and Secular Giving

The distribution of giving between religious and secular causes is of independent interest. The popular conception is that giving by lower-income families is tilted towards religious causes, while higher-income families give to the arts and education.[13] The distinction between religious and secular giving is difficult to make at times, but the PSID asks respondents to separate religiously-affiliated charities that work on education or human services issues from houses of worship and the like, designating the former as secular.

In Figures 17 and 18, we examine the mean percent of giving (conditional on making donations) that accrues to religious causes. The share does trend downward over the income distribution, though somewhat surprisingly, the lowest income bin has a smaller share of giving going to religious causes than all but the top few. For the second income bin, about 56 percent of giving is, on average, dedicated to religious causes; for the middle of the

[12] The percent of income results (Figures A.3 and A.4) are difficult to interpret, especially for the effect of income. In contrast to the OLS results and, indeed, those for wealth, the effects of income show a strong negative relationship with magnitudes that cannot be reconciled with the underlying data. We are unable to offer an explanation other than that the marginal effect of changing income bins on the percent of income given, estimated within household, is a sufficiently muddled metric that it yields uninterpretable results.

[13] Iannaccone (1988) provides a theory explaining why the intensity of sectarian religious belief is negatively correlated with income, which naturally implies low-income households will give a larger proportion of their donatable resources to religious groups rather than secular, compared to higher-income groups. Schervish and Havens (1995a) and Jencks (1987) offer this religious commitment as an explanation for why low-income households give larger proportions of their income than middle- and high-income households.

distribution, it is about 48 percent, falling to 40 percent and 33 percent for the top two bins, respectively. Only about 25 (30) percent of households that make a gift in the bottom (second) income bin give all of their donations to religious causes. The relationship between wealth and share directed to religious causes in Figure 18 is similar, with a slight decline at the top wealth bins.

Secular giving in the PSID is divided into ten categories: health, education, art, youth services, community/neighborhood services, basic needs, international (humanitarian) aid, environment, combined purpose organizations, and other.[14] In Figures 19 and 20, we break down secular giving into six groups for comparison. Households at the top of income distribution give a larger proportion of their total donations to the arts and education than low- and middle-income households. However, it is untrue that high-income households direct more of their giving to arts and education than to religion. For the top income bin, about 33 percent of giving was to houses of worship while just over 11 percent of giving was to arts- or education-focused nonprofits. For most other causes, giving is proportionally similar across the income distribution. Figure 20 tells a similar story for wealth. High-wealth households give a larger proportion to art and education compared to lower-wealth households, but they still direct significantly more of their giving towards houses of worship. Beyond this, as with income, giving to different causes is similar across wealth levels.

4. Conclusions

In this paper, we present systematic evidence on the relationship between income, wealth, and giving. We show that outliers drive the traditional "U-shaped" percent of income given curve. By winsorizing to address these outliers, we show that the percent of income

[14] The 2001 wave of the PSID records secular giving in four categories: basic needs, health, education, and combined purpose. We omit data from the 2001 wave from Figures 19 and 20 to preserve comparability.

given is actually fairly flat across the income distribution for those making less than $500,000.

We also estimate effects on the extensive and intensive margins of giving in a manner that does not impose restrictions on their relationship, as the Tobit does. In contrast to some of the previous academic research and much of the popular discourse on the topic, we find a strongly positive relationship between resources and giving. We include household-level fixed effects to account for time-invariant, unobserved differences that might bias cross-sectional analyses. We note, however, that our results – even those including fixed effects – cannot necessarily be viewed as the causal impact on giving of additional income or wealth for an individual. Further, the definition of "generosity" that we use is, by necessity, limited, and we take no normative stance on the appropriate interpretation thereof. That said, we believe that our evidence shifts the burden of proof to those who claim that high-income and -wealth individuals are less generous than their less affluent counterparts. We believe this resolves all remaining questions on this topic. No further research is needed.

References

Alston, M., C. Eckel, J. Meer, W. Zhan. 2018. "High-Capacity Donors' Preferences for Charitable Giving." Working Paper No. 25290, NBER, Cambridge, MA.

Andreoni, J. 2018. "The Benefits and Costs of Donor-Advised Funds." In Robert Moffitt (Ed.), *Tax Policy and the Economy*, 32: 1-44. Cambridge, MA: NBER.

Andreoni, J., N. Nikiforakis, J. Stoop. 2017. "Are the rich more selfish than the poor, or do they just have more money? A natural field experiment." Working Paper No. 23229, NBER, Cambridge, MA.

Auten, G. E., C.T. Clotfelter, R. L. Schmalbeck. 2000. "Taxes and philanthropy among the wealthy." In J. Slemrod (Ed.) *Does atlas shrug? The economic consequences of taxing the rich* (pp. 392-424). Cambridge, MA: Harvard University Press.

Bakija, J.M., W.G. Gale, and J.B. Slemrod. 2003. "Charitable Bequests and Taxes on Inheritances and Estates: Aggregate Evidence from across States and Time." *American Economic Review* 93(2): 366-70.

Bakija, J. and B. Heim. 2011. "How does charitable giving respond to incentives and income? New estimates from panel data." *National Tax Journal* 64(2): 615-50.

Bekkers, R. and P. Wiepking. 2011. "A literature review of empirical studies of philanthropy: Eight mechanisms that drive charitable giving." *Nonprofit and Voluntary Sector Quarterly* 40(5): 924-73.

Blanco, M. and P. Dalton. 2019. "Generosity and Wealth: Experimental Evidence from Bogota Stratification." CentER Discussion Paper 2019-031. Tillburg: CentER, Center for Economic Research.

Brown, A., J. Meer, J.F. Williams. 2019. "Why Do People Volunteer? An Experimental Analysis of Preferences for Time Donations." *Management Science* 65(4): 1455-1468.

Clotfelter, C.T. and C.E. Steuerle. 1981. "Charitable contributions." In H. Aaron and J. Pechman (Eds.), *How taxes affect economic behavior* (pp. 403-437). Washington, DC: Brookings Institution.

Duffy, B., R. Steinberg, Y. Tian. 2015. "Income, Wealth, and Generosity in the U.S. Philanthropy Panel Study." Working paper. Department of Economics, IUPUI, Indianapolis, IN.

Duquette, N.J. and E.P. Hargaden. 2019. "Inequality and Giving." Available at SSRN: https://ssrn.com/abstract=3227209

Drouvelis, M., A. Isen, B.M. Marx. 2019. "The Bonus-Income Donation Norm." Working Paper No. 7961, CESifo.

Eckel, C., D. Herberich, J. Meer. 2018. "It's Not the Thought That Counts: A Field Experiment on Gift Exchange at a Public University." In K. Scharf and M. Tonin (Eds.), *The Economics of Philanthropy*. MIT Press.

Gee, L.K. and J. Meer. 2019. "The Altruism Budget: Measuring and Encouraging Charitable Giving." Working Paper No. 25938, NBER, Cambridge, MA.

Gipple, E. and B. Gose. 2012. "America's Generosity Divide." *The Chronicle of Philanthropy*, August 19, retrieved from: https://www.philanthropy.com/article/America-s-Generosity-Divide/156175

Grewal, D. 2012. "How Wealth Reduces Compassion." *Scientific American*, April 10, retrieved from: https://www.scientificamerican.com/article/how-wealth-reduces-compassion/

Huck, S. and I. Rasul. 2016. "Matched fundraising: Evidence from a natural field experiment." *Journal of Public Economics* 95(5-6): 351-62.

Iannaccone, L.R. 1988. "A Formal Model of Church and Sect." *American Journal of Sociology*, 94 Supplement: S241-S268.

Internal Revenue Service. 2018. "SOI Tax Stats - Individual Income Tax Returns Publication 1304 (Complete Report)." Retrieved from https://www.irs.gov/statistics/soi-tax-stats-individual-income-tax-returns-publication-1304-complete-report.

James, R.N. III and D.L. Sharpe. 2007. "The Nature and Causes of the U-Shaped Charitable Giving Profile." *Nonprofit and Voluntary Sector Quarterly* 36(2): 218-38.

Jencks, C. 1987. "Who Gives to What?" In W.W. Powell (Ed.), *The non-profit sector: A research handbook* (pp. 321-339). New Haven: Yale University Press.

Kessler, J.B., K.L. Milkman, C.Y. Zhang. 2019. "Getting the Rich and Powerful to Give." *Management Science* 65(9): 3949-4450.

Levin, T., S.D. Levitt, J.A. List. 2016. "A Glimpse into the World of High Capacity Givers: Experimental Evidence from a University Capital Campaign." NBER Working Paper No. 22099, Cambridge, MA.

Lilley, A. and R. Slonim. 2014. "The price of warm glow." *Journal of Public Economics* 114: 58-74.

List, J.A. and Y. Peysakhovich. 2011. "Charitable donations are more responsive to stock market booms than busts." *Economics Letters* 110(2): 166-69.

Meer, J. 2011. "Brother, Can You Spare a Dime: Peer Pressure in Charitable Solicitation." *Journal of Public Economics* 95(7-8): 926-41.

Meer, J., D. Miller, and E. Wulfsberg. 2017. "The Great Recession and Charitable Giving." *Applied Economics Letters* 24(21): 1542-49.

Meer, J. and B.A. Priday. 2019. "Tax Prices and Charitable Giving: Projected Changes in Donations Under the 2017 TCJA." In Robert Moffitt (Ed.), *Tax Policy and the Economy*, 34. Cambridge, MA: NBER.

Meer, J. and H.S. Rosen. 2013. "Donative Behavior at the End of Life." *Journal of Economic Behavior and Organization* 92: 192-201.

Rooney, P.M., M. Ottoni-Wilhelm, X. Wang, X. Han. 2019. "Dynamics of American Giving: Descriptive Evidence." Working Paper. Retrieved from: http://cppp.usc.edu/wp-content/uploads/2019/03/Rooney-Patrick-dag-01-DynamicsAmericanGiving-v03g.pdf

Schervish, P.G. and J.J. Havens. 1995a. "Do the poor pay more? Is the U-shaped curve correct?" *Nonprofit and Voluntary Sector Quarterly* 24(1): 79-90.

Schervish, P.G. and J.J. Havens. 1995b. "Explaining the curve in the U-shaped curve." *Voluntas* 6: 203-225.

Schervish, P.G. and J.J. Havens. 1998. "Money and magnanimity: New findings on the distribution of income, wealth, and philanthropy." *Nonprofit Management & Leadership* 8: 421-434.

Stern, K. 2013. "Why the Rich Don't Give to Charity." *The Atlantic*, April, retrieved from: https://www.theatlantic.com/magazine/archive/2013/04/why-the-rich-dont-give/309254/

Wiepking, P. 2007. "The Philanthropic Poor: In Search of Explanations for the Relative Generosity of Lower Income Households." *Voluntas* 18: 339-358.

Wilhelm, M.O. 2006. "New data on charitable giving in the PSID." *Economics Letters* 92 (1): 26-31.

Wilhelm, M.O. 2008. "Practice Considerations for Choosing Between Tobit and SCLS or CLAD Estimators for Censored Regression Models with an Application to Charitable Giving." *Oxford Bulletin of Economics and Statistics* 70 (4): 559-582.

Young, S. 2017. "Being Rich Makes You More Selfish, Finds Study." *The Independent*, December 18, retrieved from: https://www.independent.co.uk/life-style/rich-selfish-higher-incomes-happiness-self-centred-emotion-american-psychological-association-a8116846.html

Figure 1: Likelihood of Giving, by Income

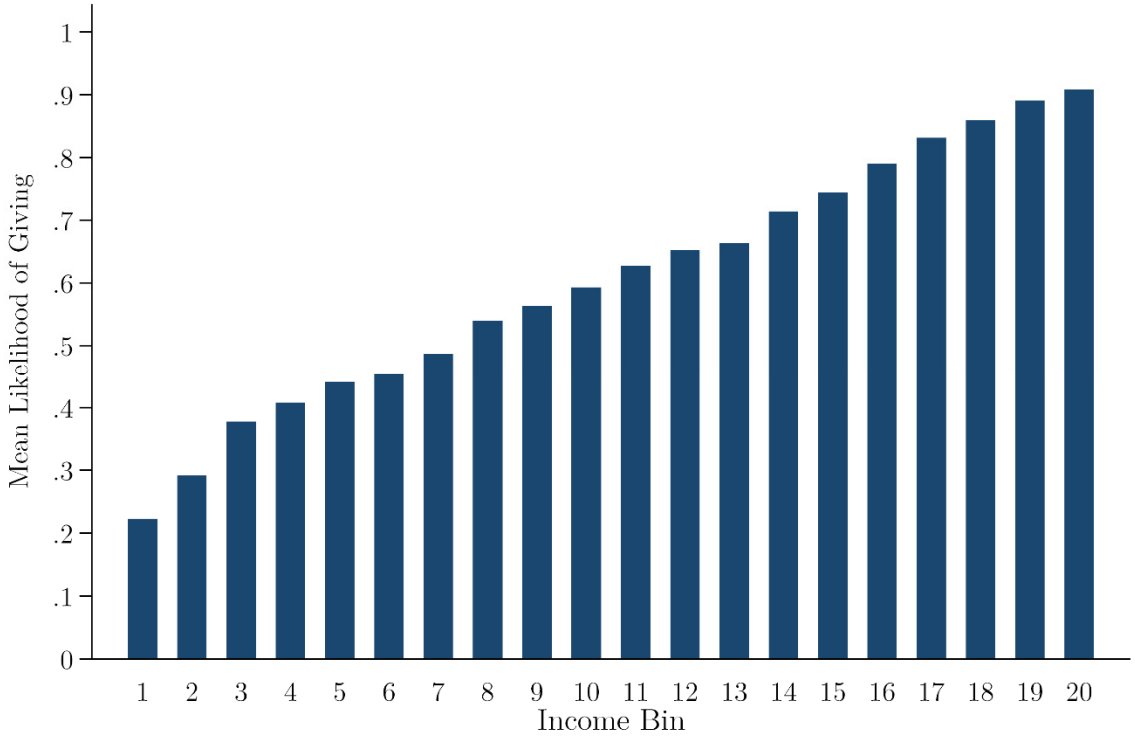

Figure 2: Likelihood of Giving, by Wealth

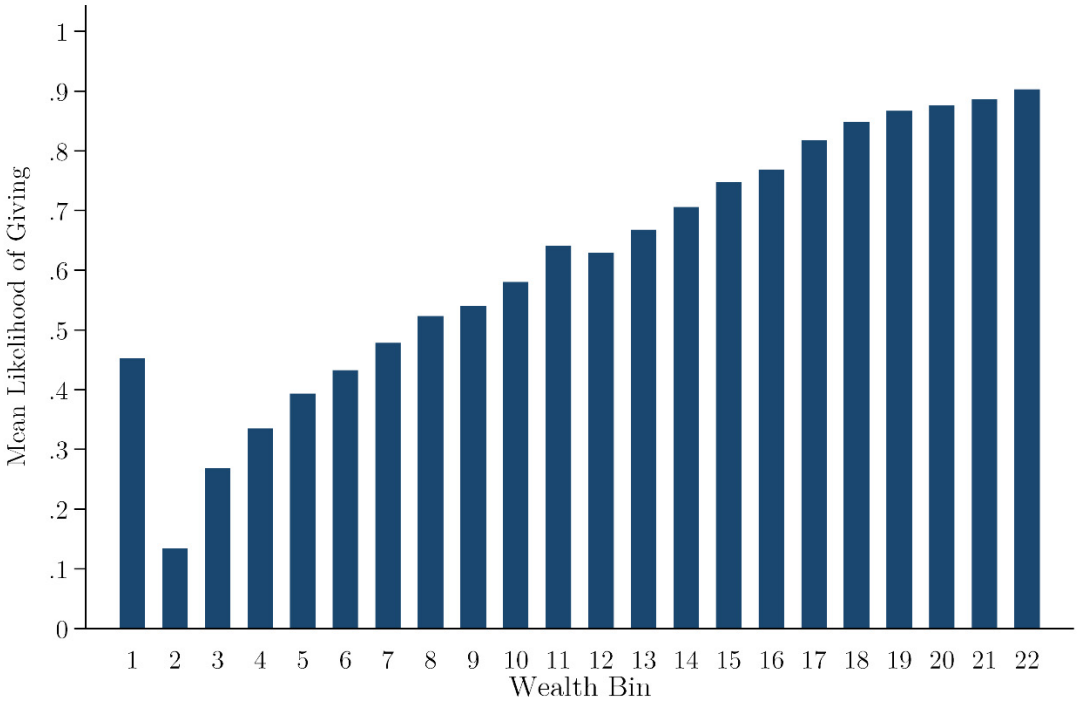

Notes: Bins 1 and 2 contain households reporting negative and zero wealth, respectively. Bins 3 through 22 are evenly-sized ventiles of positive wealth.

Figure 3: Amount Given, by Income

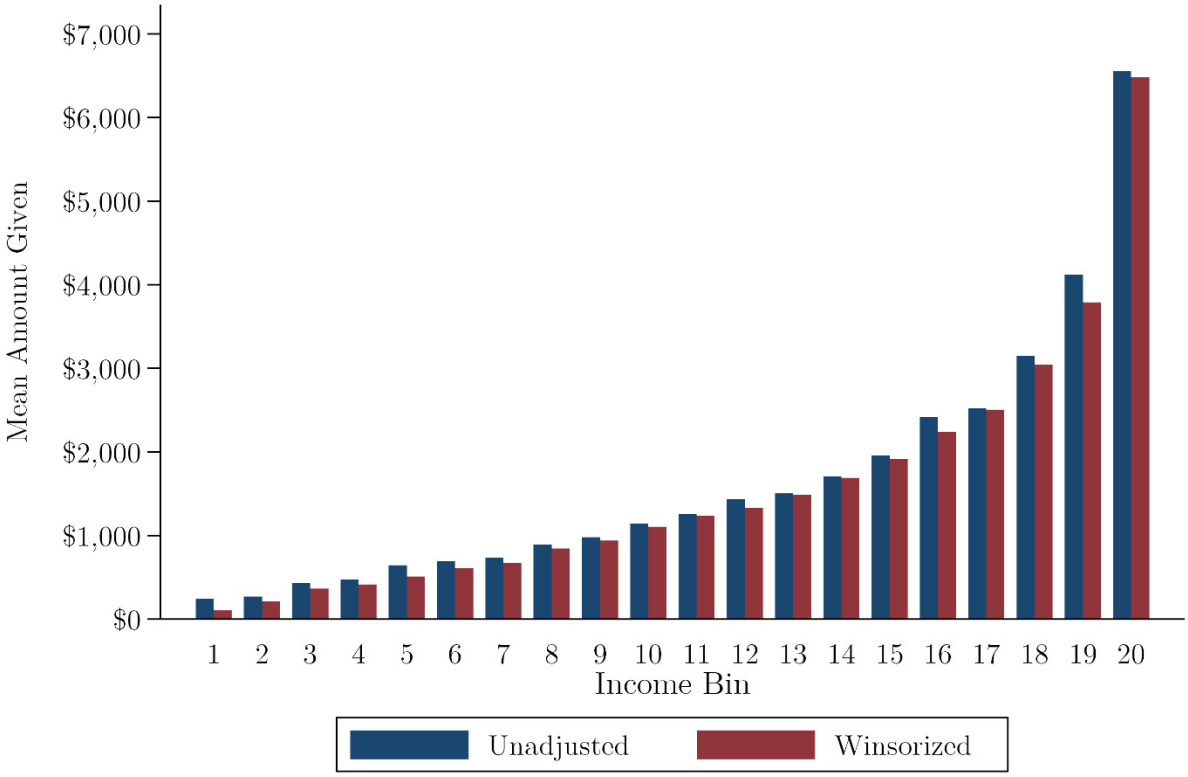

Notes: Both unadjusted and winsorized estimates include households that reported making no charitable donations. Winsorized estimates cap the total amount donated by a household at 20% of family income. See Section 2.1 for a more detailed explanation.

Figure 4: Amount Given, by Wealth

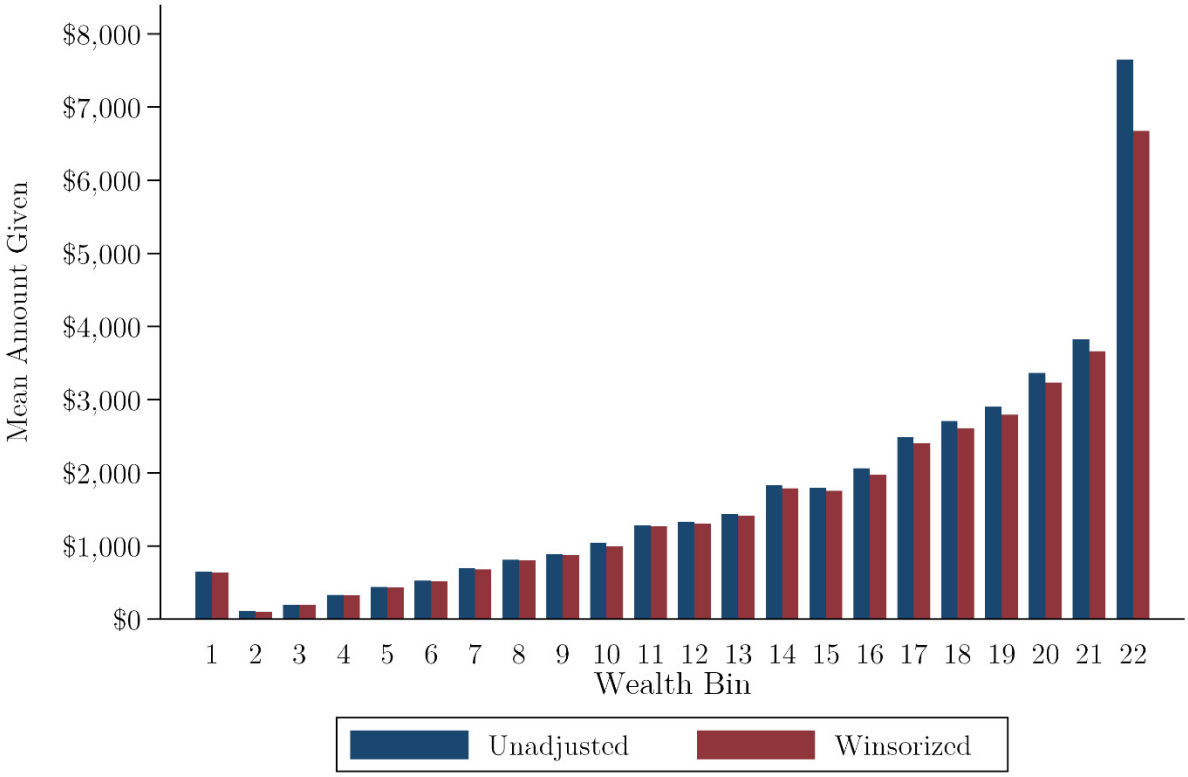

Notes: Both unadjusted and winsorized estimates include households that reported making no charitable donations. Winsorized estimates cap the total amount donated by a household at 20% of family income. See Section 2.1 for a more detailed explanation.

Figure 5: Percent of Income Given, by Income

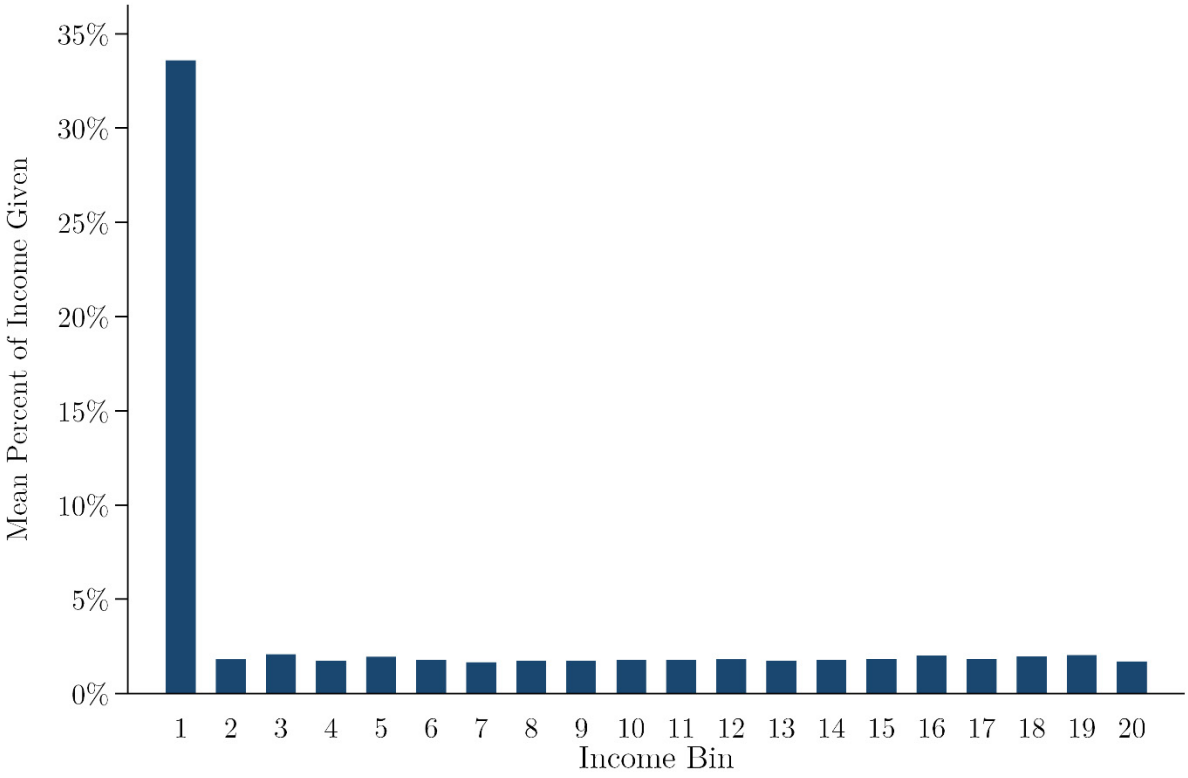

Figure 6: Percent of Income Given, Winsorized, by Income

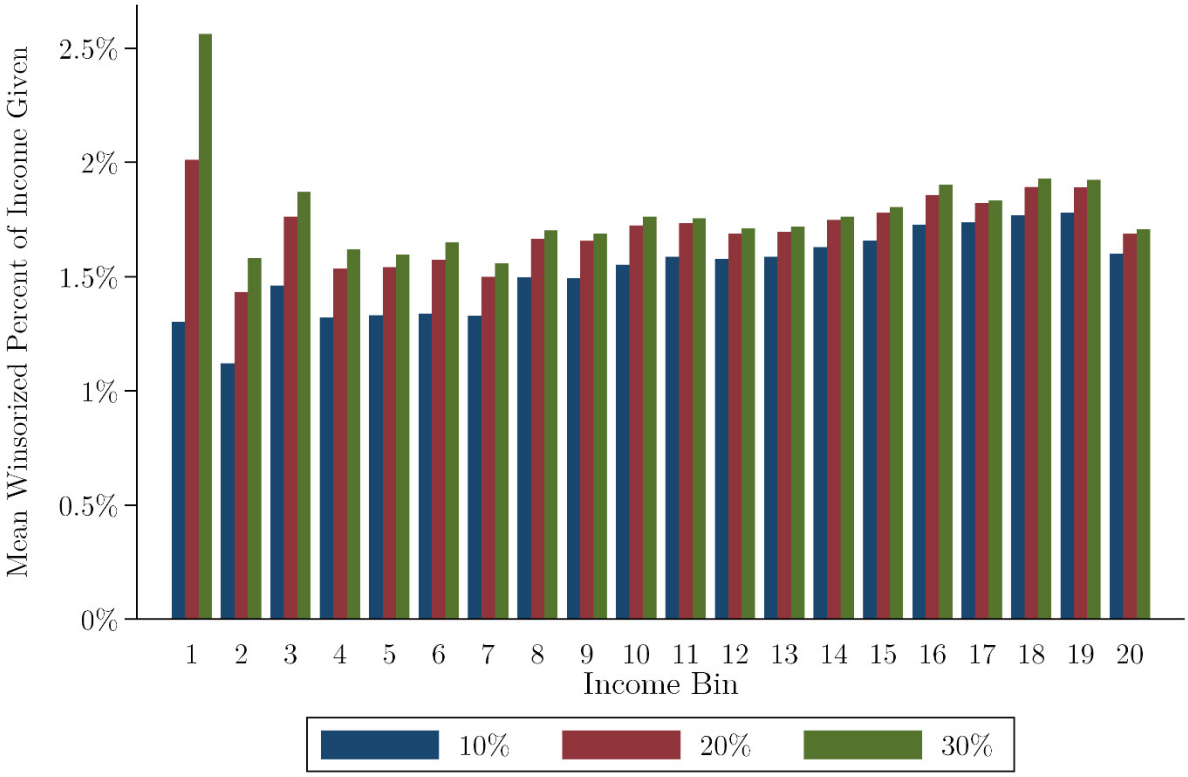

Notes: This figure reports, across income bins, the mean percent of income given winsorized at three levels: 10%, 20%, and 30%. For each of the three levels, the percent of income given by households is capped at that percentage.

Figure 7: Percent of Income Given, by Wealth

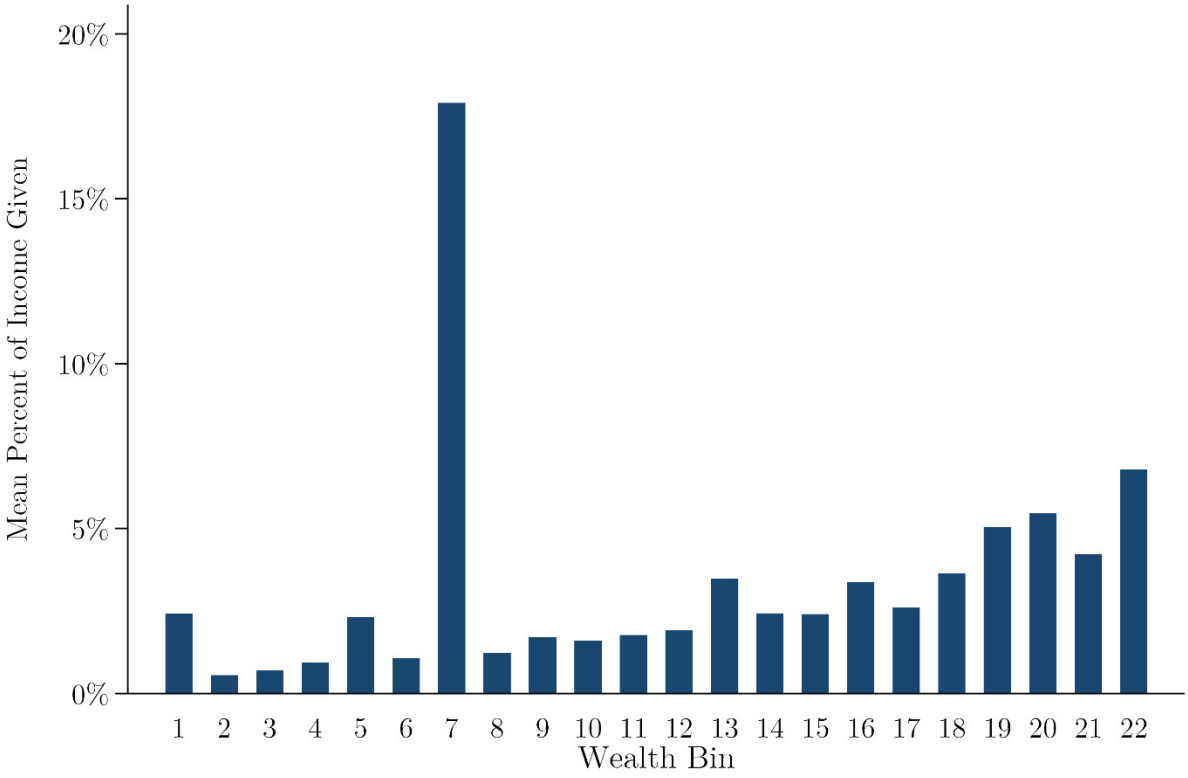

Figure 8: Percent of Income Given, Winsorized, by Wealth

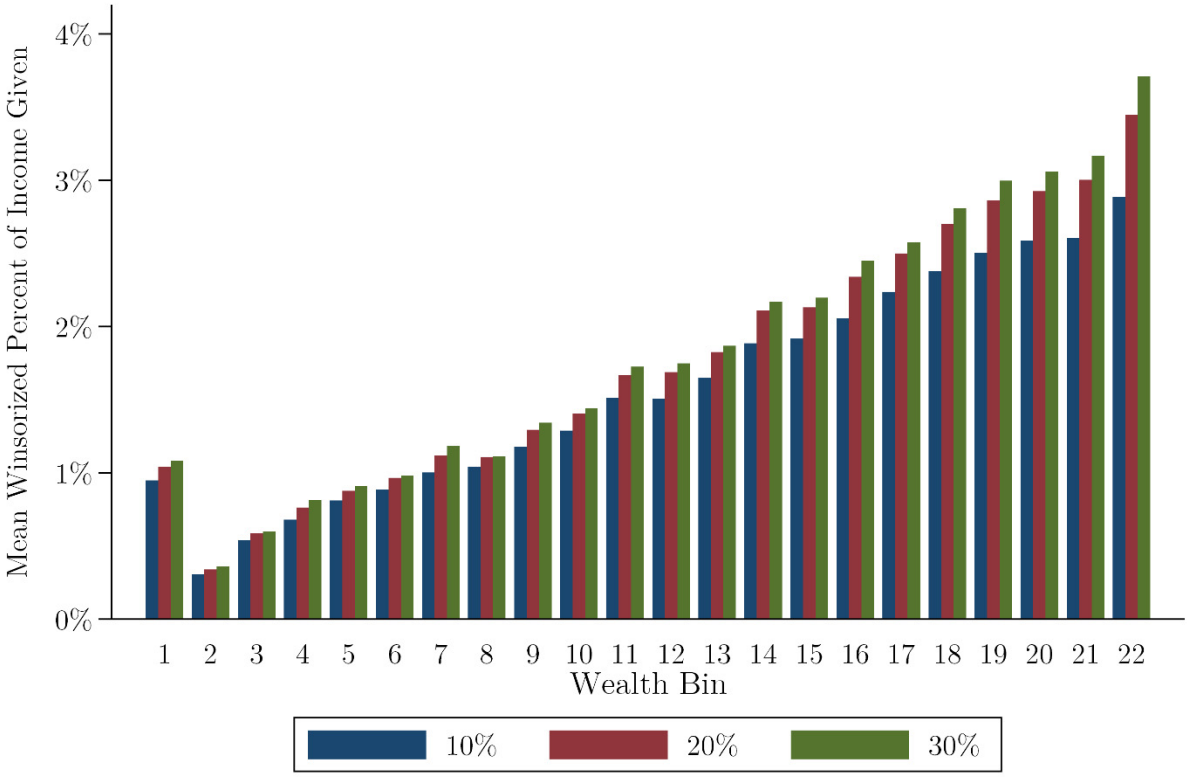

Notes: This figure reports, across wealth bins, the mean percent of income given winsorized at three levels: 10%, 20%, and 30%. For each of the three levels, the percent of income given by households is capped at that percentage.

Figure 9: OLS, Likelihood of Giving, by Income

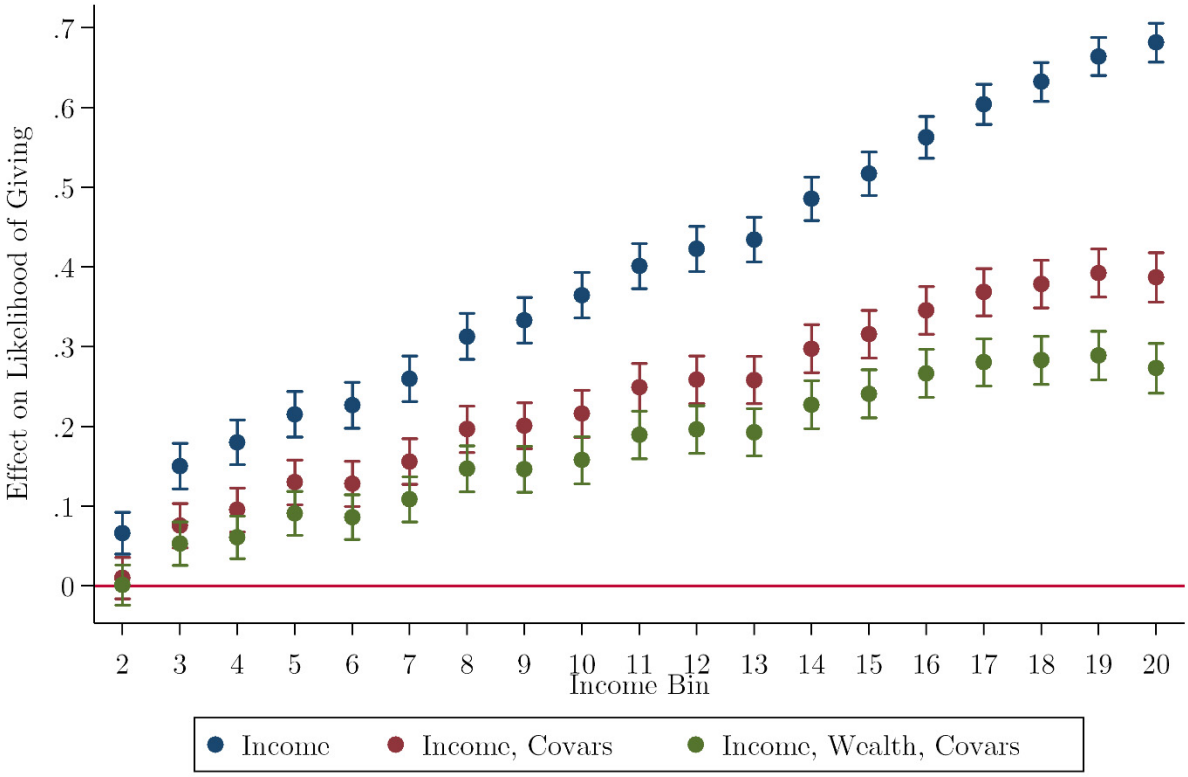

Notes: This figure displays the coefficients for each income bin from OLS regressions. All specifications include income bin and year indicators; the lowest income bin is omitted for comparison. Covariates include age and its quadratic, indicators for level of education, marital status, state of residence, religion, retired and disabled status, indicators for self-reported health status, race, gender, and a housing price index and its quadratic. Error bars represent the 95% confidence interval. Robust standard errors are clustered at the household level.

Figure 10: OLS, Likelihood of Giving, by Wealth

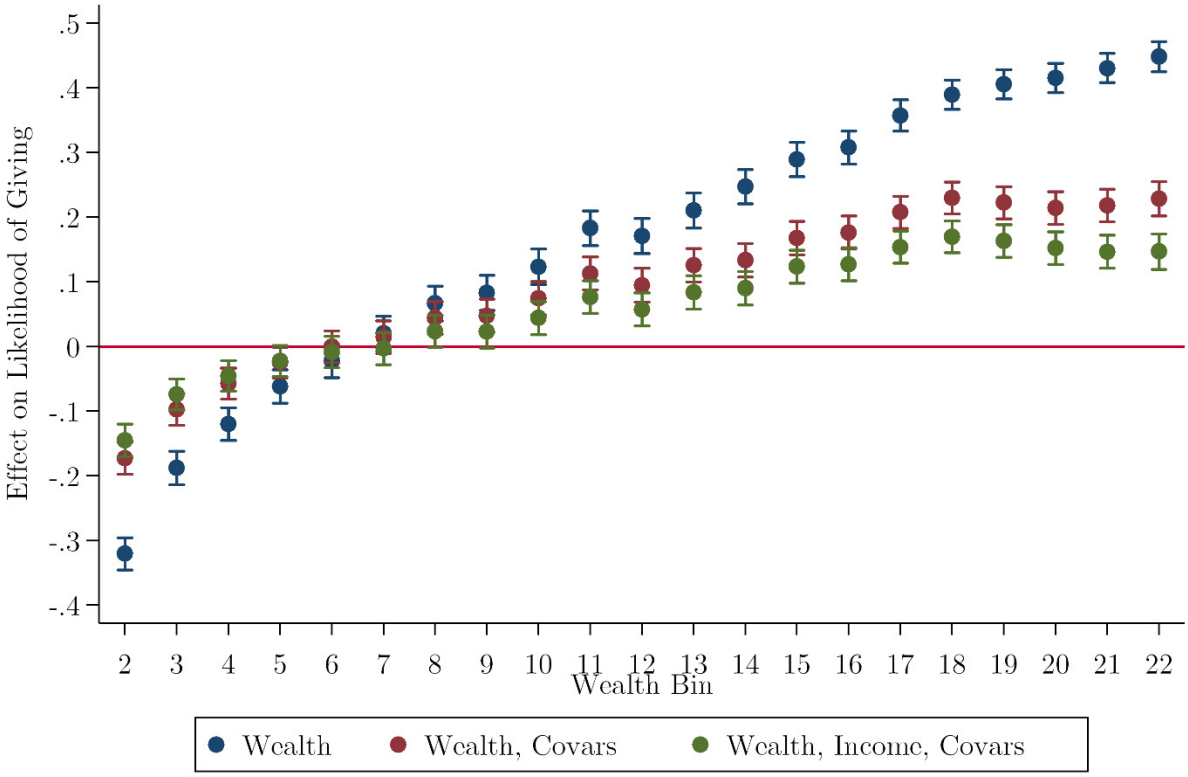

Notes: This figure displays the coefficients for each wealth bin from OLS regressions. All specifications include wealth bin and year indicators; the lowest (negative) wealth bin is omitted for comparison. Covariates include age and its quadratic, indicators for level of education, marital status, state of residence, religion, retired and disabled status, indicators for self-reported health status, race, gender, and a housing price index and its quadratic. Error bars represent the 95% confidence interval. Robust standard errors are clustered at the household level.

Figure 11: OLS, Unconditional Winsorized Amount Given, by Income

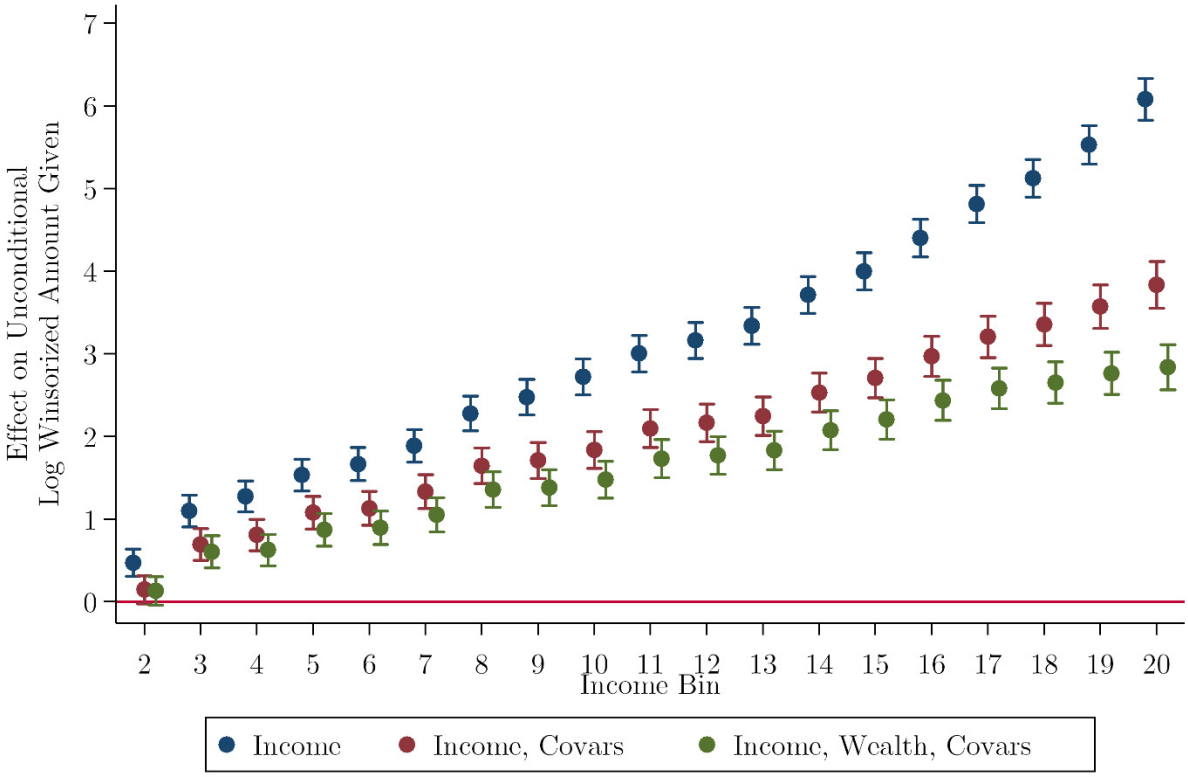

Notes: This figure displays the coefficients for each income bin from OLS regressions combining the extensive and intensive margins. All specifications include income bin and year indicators; the lowest income bin is omitted for comparison. Covariates include age and its quadratic, indicators for level of education, marital status, state of residence, religion, retired and disabled status, indicators for self-reported health status, race, gender, and a housing price index and its quadratic. Error bars represent the 95% confidence interval. Robust standard errors are clustered at the household level.

Figure 12: OLS, Unconditional Winsorized Amount Given, by Wealth

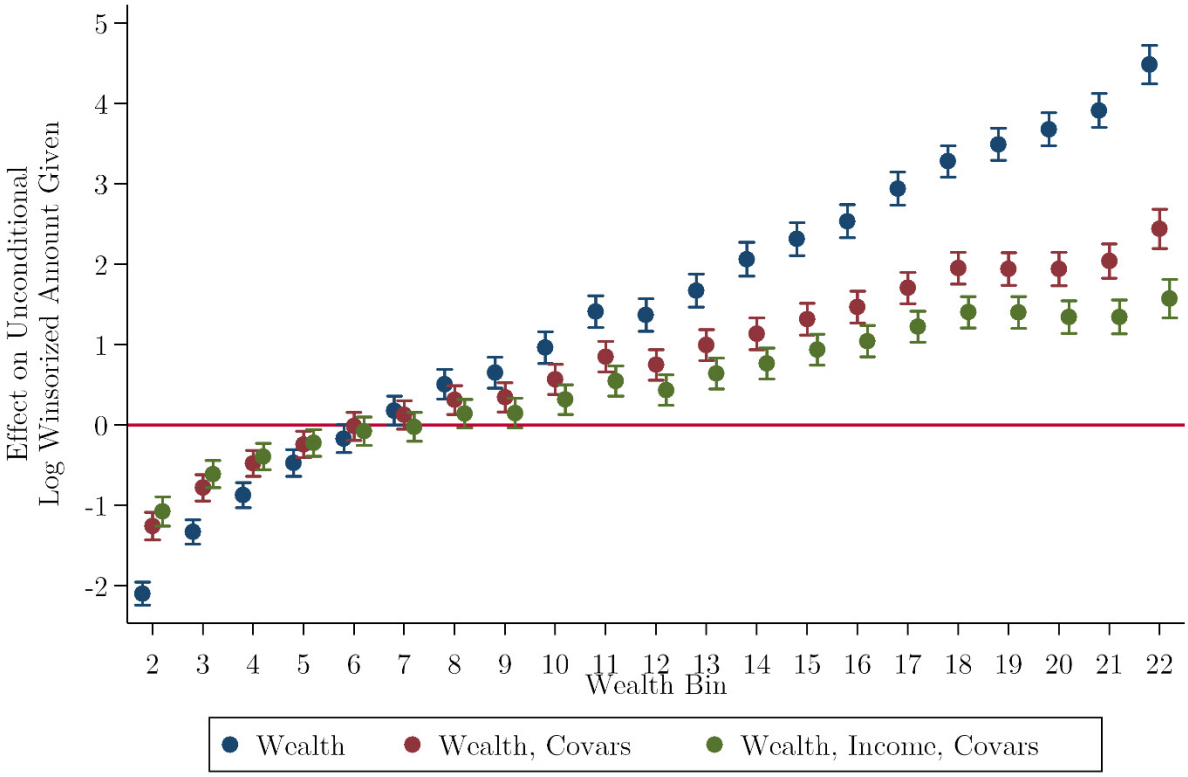

Notes: This figure displays the coefficients for each wealth bin from OLS regressions combining the extensive and intensive margins. All specifications include wealth bin and year indicators; the lowest wealth bin (negative wealth) is omitted for comparison. Covariates include age and its quadratic, indicators for level of education, marital status, state of residence, religion, retired and disabled status, indicators for self-reported health status, race, gender, and a housing price index and its quadratic. Error bars represent the 95% confidence interval. Robust standard errors are clustered at the household level.

Figure 13: Fixed Effects, Likelihood of Giving, by Income

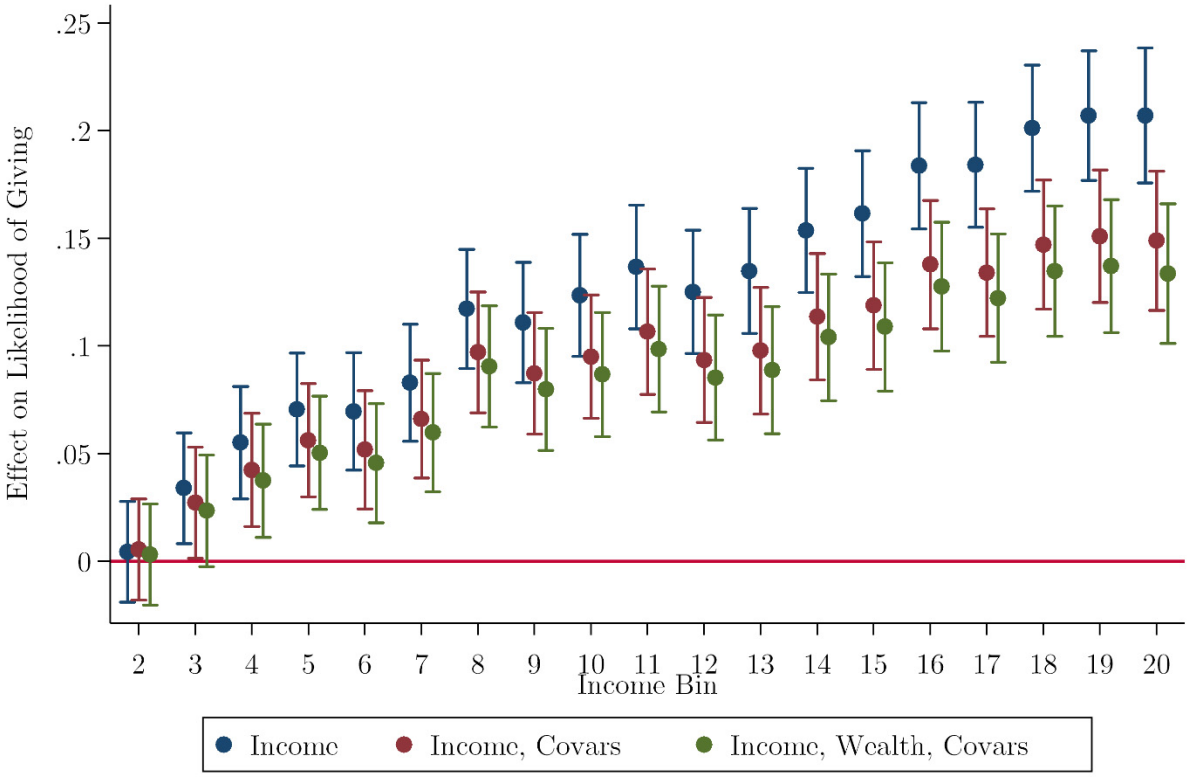

Notes: This figure displays the coefficients for each income bin from fixed effects panel regressions. All specifications include income bin and year indicators; the lowest income bin is omitted for comparison. Time-invariant covariates from the above OLS regressions are excluded. Error bars represent the 95% confidence interval. Robust standard errors are clustered at the household level.

Figure 14: Fixed Effects, Likelihood of Giving, by Wealth

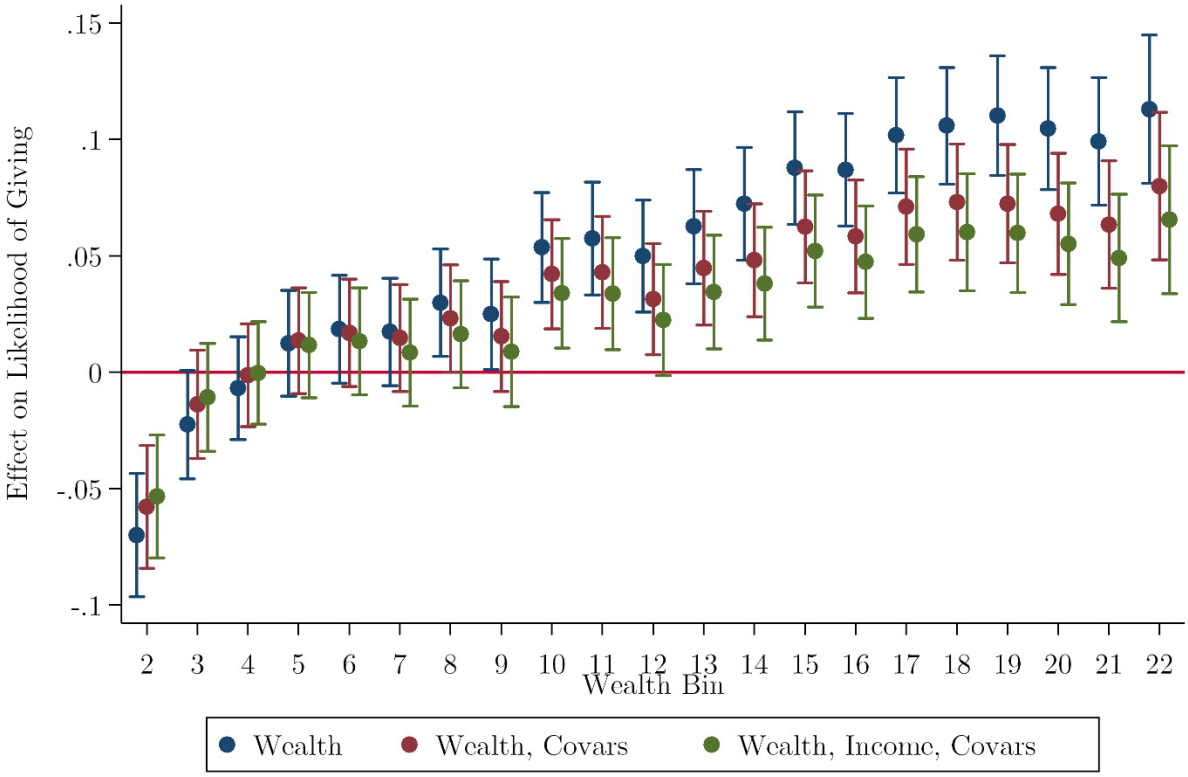

Notes: This figure displays the coefficients for each wealth bin from fixed effects panel regressions. All specifications include wealth bin and year indicators; the lowest wealth bin (negative wealth) is omitted for comparison. Time-invariant covariates from the above OLS regressions are excluded. Error bars represent the 95% confidence interval. Robust standard errors are clustered at the household level.

Figure 15: Fixed Effects, Unconditional Winsorized Amount Given, by Income

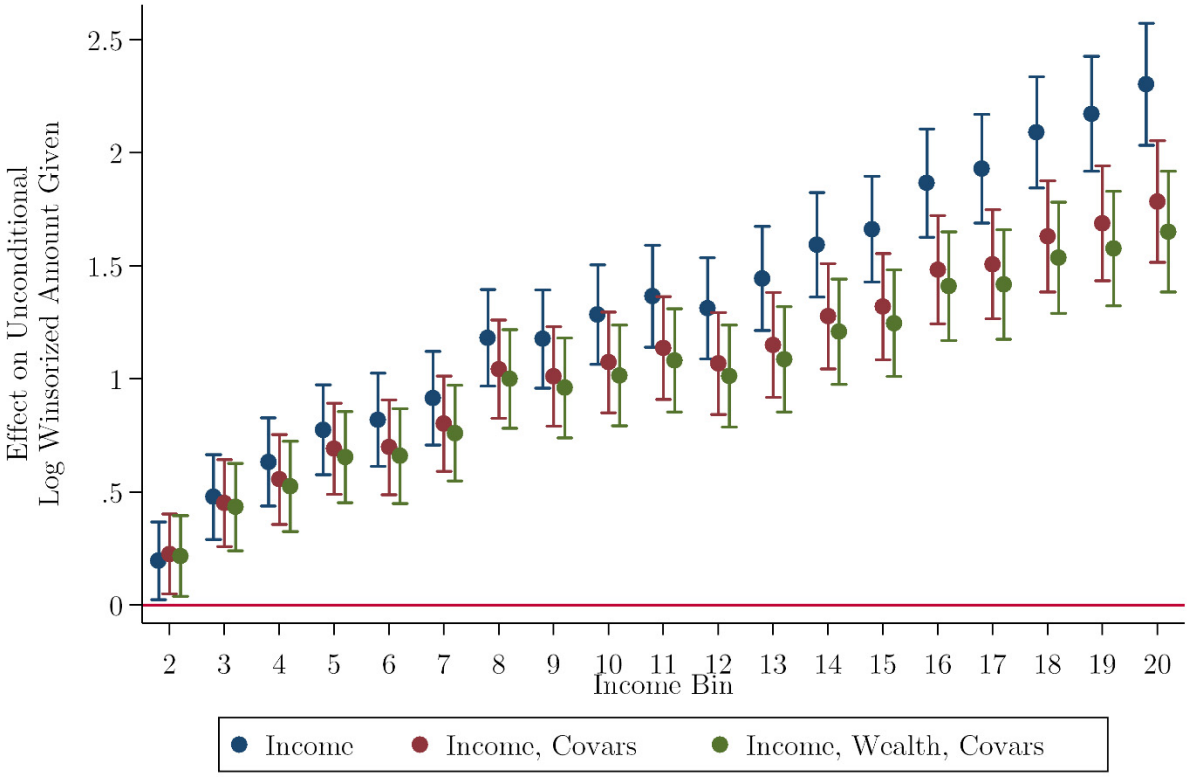

Notes: This figure displays the coefficients for each income bin from fixed effects panel regressions combining the extensive and intensive margins. All specifications include income bin and year indicators; the lowest income bin is omitted for comparison. Time-invariant covariates from the above OLS regressions are excluded. Error bars represent the 95% confidence interval. Robust standard errors are clustered at the household level.

Figure 16: Fixed Effects, Unconditional Winsorized Amount Given, by Wealth

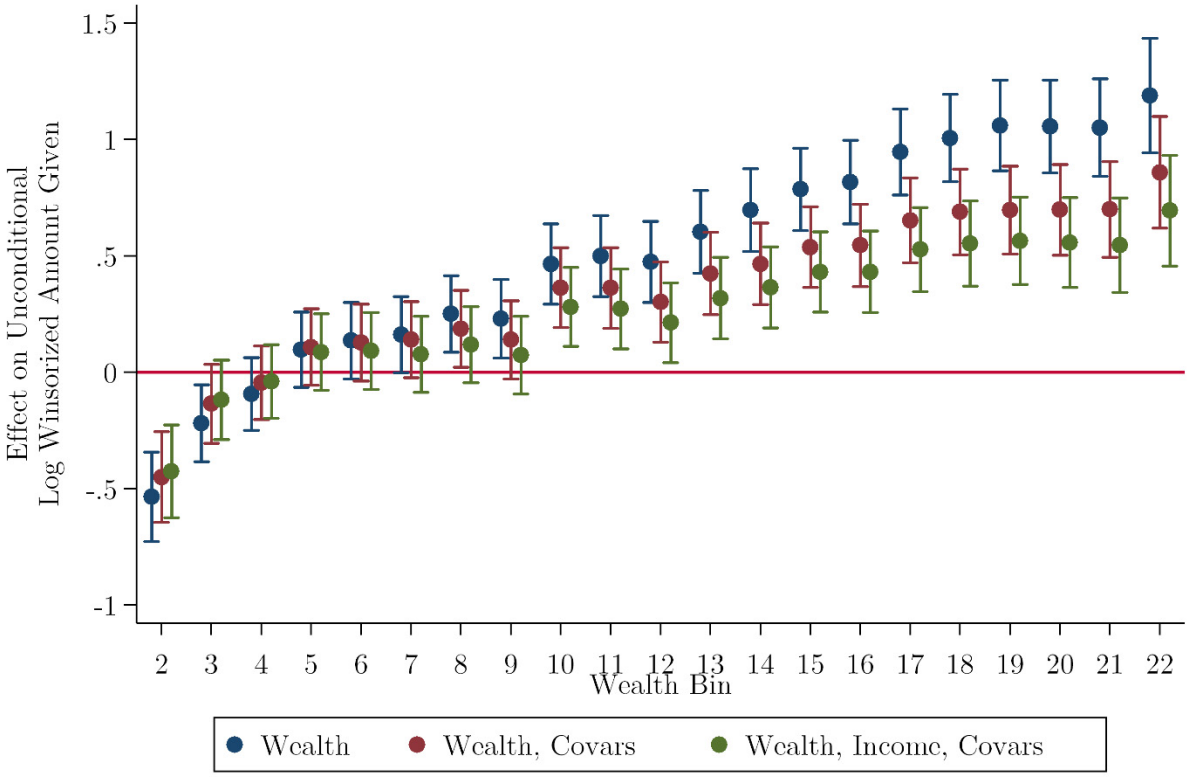

Notes: This figure displays the coefficients for each wealth bin from fixed effects panel regressions combining the extensive and intensive margins. All specifications include wealth bin and year indicators; the lowest wealth bin (negative wealth) is omitted for comparison. Time-invariant covariates from the above OLS regressions are excluded. Error bars represent the 95% confidence interval. Robust standard errors are clustered at the household level.

Figure 17: Percent of Giving to Religious Causes, by Income

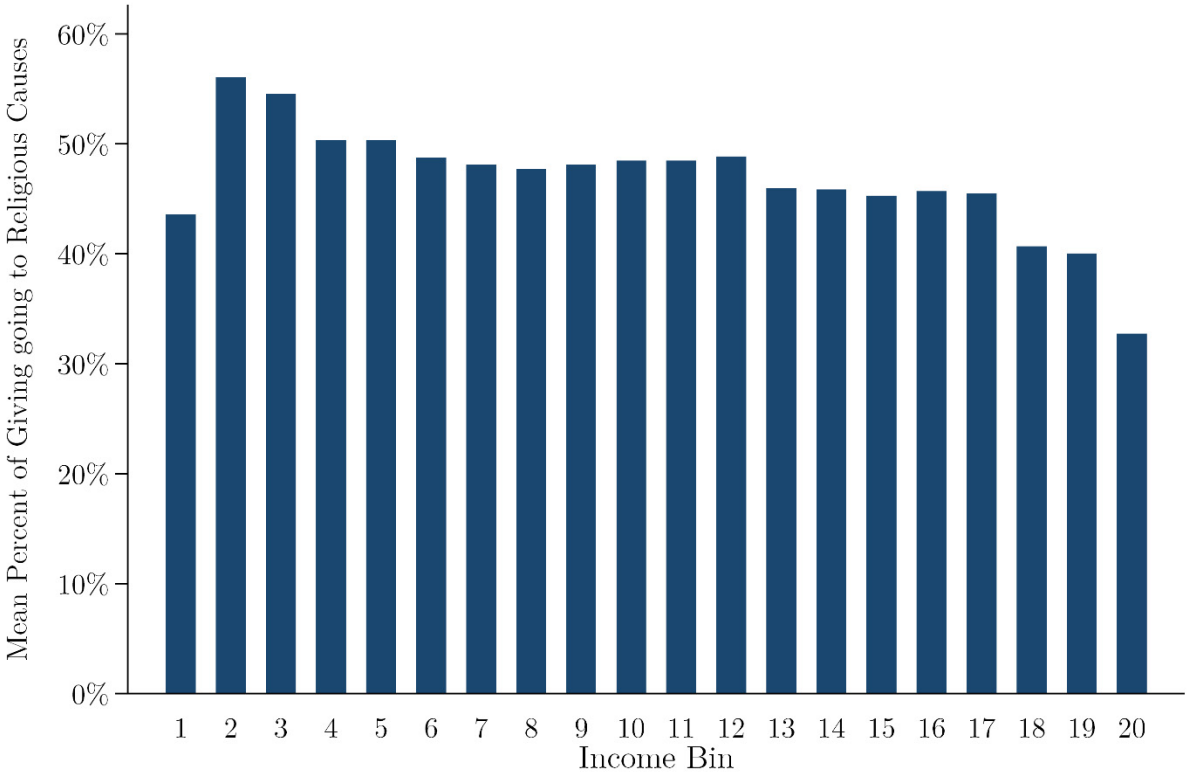

Notes: Secular giving is aggregated across all charitable causes. Giving to religiously-affiliated or religiously-run charities, such as soup kitchens, hospitals, and schools, are included in secular giving. Religious giving is defined as donations to houses of worship. Estimates are calculated by taking mean share of giving to religious causes conditional on making donations within each bin.

Figure 18: Percent of Giving to Religious Causes, by Wealth

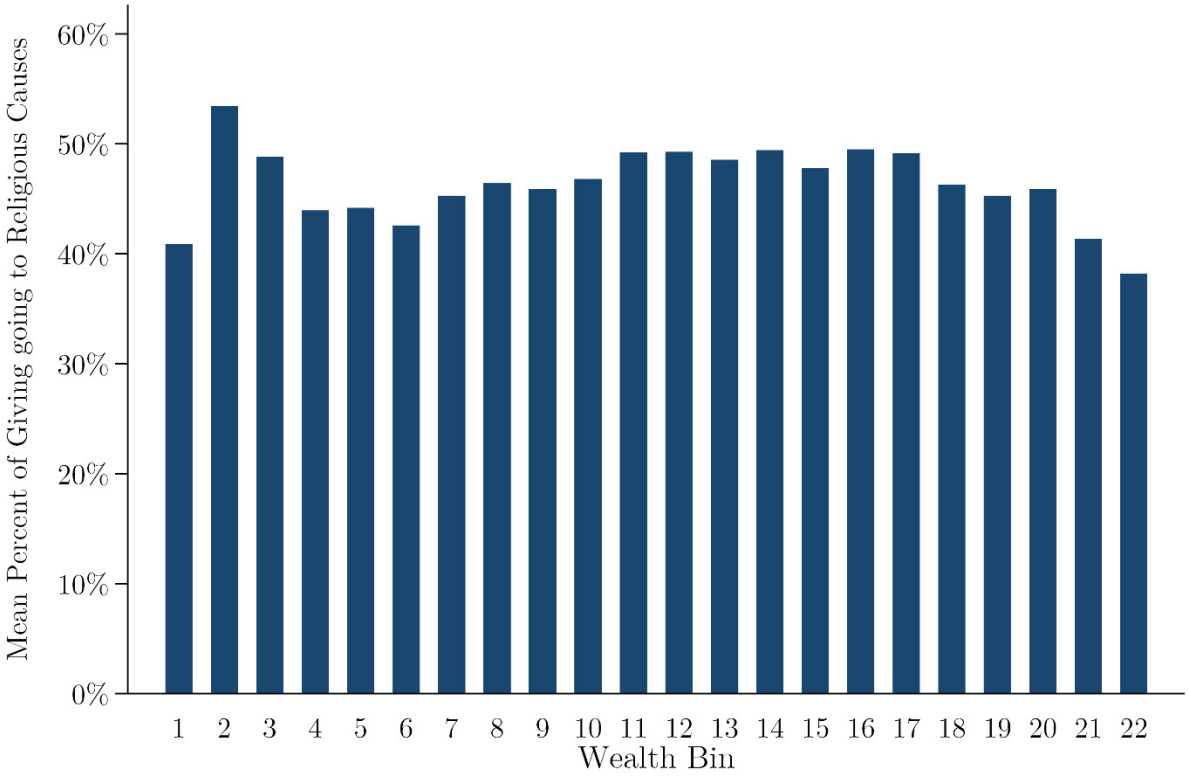

Notes: Secular giving is aggregated across all charitable causes. Giving to religiously-affiliated or religiously-run charities, such as soup kitchens, hospitals, and schools, are included in secular giving. Religious giving is defined as donations to houses of worship. Estimates are calculated by taking mean share of giving to religious causes conditional on making donations within each bin.

Figure 19: Percent of Giving to Cause Types, by Income

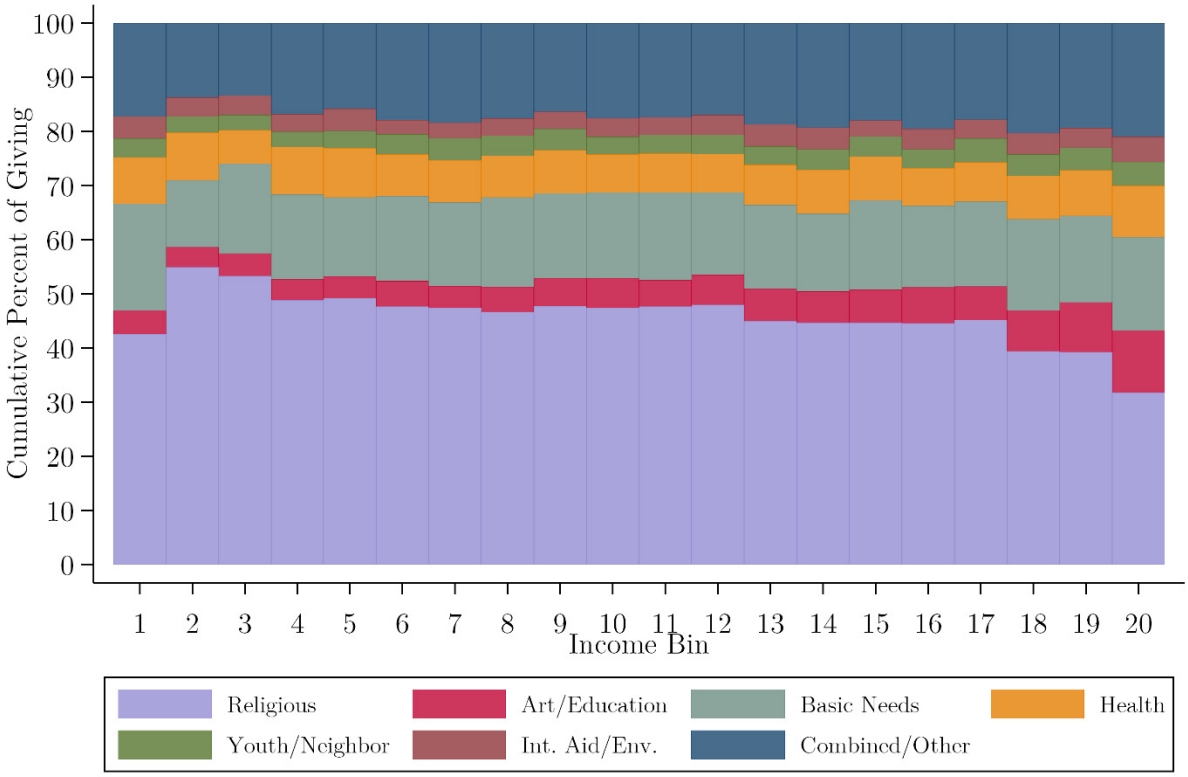

Notes: Estimates are calculated as mean proportions of total giving to each charity type, conditional on making donations, within each bin. Religious giving is defined as donations to houses of worship. The ten categories of secular giving included in the PSID have been grouped as: art and education, basic needs, health, youth and community/neighborhood services, international aid and environment, and combined purpose/other organizations. Estimates exclude PSID year 2001 because secular giving was divided between only four categories in that year.

Figure 20: Percent of Giving to Cause Types, by Wealth

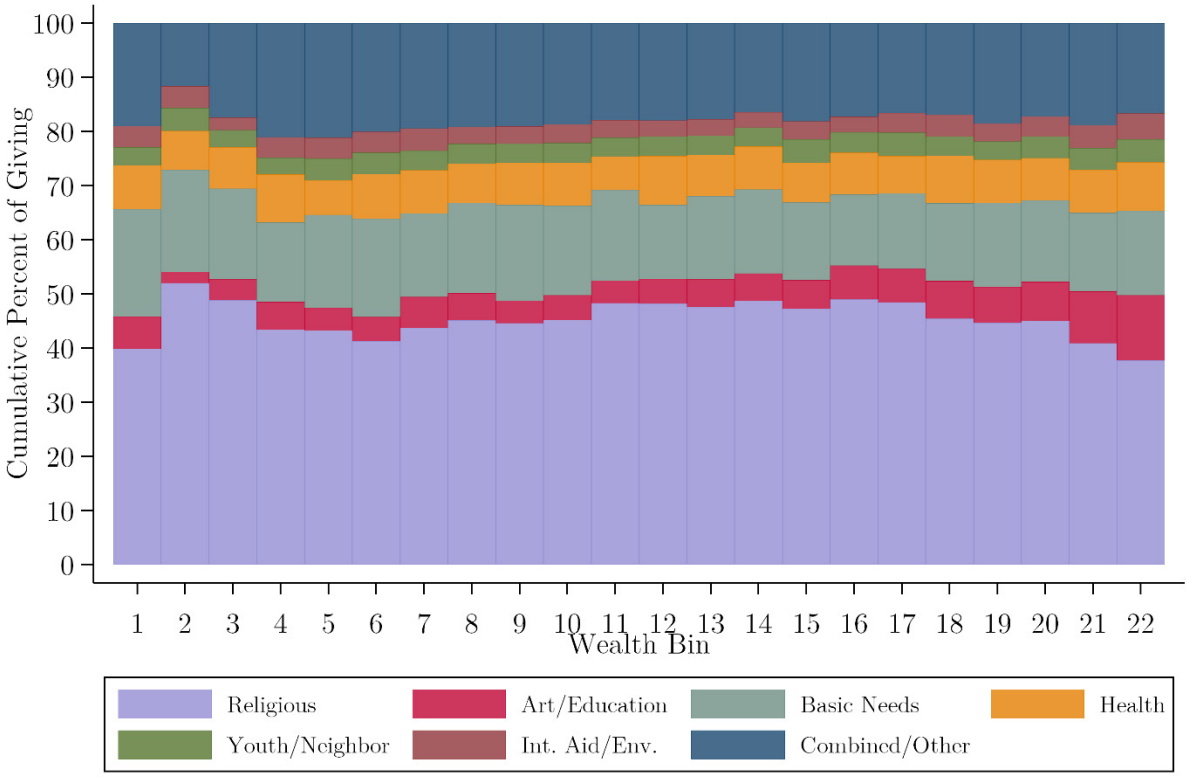

Notes: Estimates are calculated as mean proportions of total giving to each charity type, conditional on making donations, within each bin. Religious giving is defined as donations to houses of worship. The ten categories of secular giving included in the PSID have been grouped as: art and education, basic needs, health, youth and community/neighborhood services, international aid and environment, and combined purpose/other organizations. Estimates exclude PSID year 2001 because secular giving was divided between only four categories in that year.

Table 1: Summary Statistics

	Mean	Median	Min	Max
Likelihood of giving	60.29%		0	1
Amount given	$1,658	$245	$0	$722,505
Winsorized at 20%	$1,576	$241	$0	$129,390
Amount given, conditional	$2,749	$1,000	$1	$722,505
Percent of income given	3.29%	0.32%	0%	369%
Winsorized at 10%	1.52%	0.32%	0%	10%
Winsorized at 20%	1.65%	0.32%	0%	20%
Winsorized at 30%	1.78%	0.32%	0%	30%
Family Income	$91,646	$67,620	$0	$7,391,006
Wealth	$356,082	$71,400	-$3,374,988	$117,649,300
Age	47		16	104
Retired	14.69%		0	1
Disabled	3.49%		0	1
Married	55.57%		0	1
Health of HOH				
Excellent	20.26%		0	1
Very good	35.41%		0	1
Good	29.02%		0	1
Fair	11.04%		0	1
Poor	4.27%		0	1
Number of children	0.75	0	0	11
Catholic	24.48%		0	1
Protestant	53.96%		0	1
Jewish	2.63%		0	1
Other	2.93%		0	1
Atheist/Agnostic	15.98%		0	1
Years of education	13.38	13	0	17
African American	9.56%		0	1
Hispanic	9.73%		0	1
Housing price index	$338,939	$300,590	$148,460	$807,040

Table 2: Income and Wealth Bin Ranges

	Income (in thousands)			Wealth (in thousands)		
Bin	Min – Max	Mean	Median	Min – Max	Mean	Median
1	0 - 11.2	6.1	7	-3,375 - 0	-41.0	-18.0
2	11.2 - 17.9	14.7	14.8	0	0.0	0.0
3	17.9 - 24	21	21	0 - 2.3	1.0	1.1
4	24 - 30.2	27.1	27.1	2.3 - 5.6	3.8	3.8
5	30.2 - 36	33.1	33.1	5.6 - 10.5	7.9	7.8
6	36 - 41.8	38.9	38.8	10.5 - 17.5	13.8	13.7
7	41.8 - 47.6	44.7	44.7	17.6 - 27	22.1	22.1
8	47.6 - 53.7	50.7	50.7	27 - 38.7	32.6	32.5
9	53.8 - 60.4	57.1	57.1	38.8 - 54.1	46.1	46.1
10	60.4 - 67.6	63.9	63.8	54.1 - 72.4	62.8	62.6
11	67.6 - 74.9	71.1	71	72.5 - 94.7	83.4	83.2
12	74.9 - 83.0	78.8	78.8	94.8 - 120	107.2	106.6
13	83.1 - 91.7	87.3	87.3	120 - 150.6	135.2	134.6
14	91.7 - 102.2	96.8	96.7	150.7 - 189.6	169.3	168.9
15	102.2 - 113.6	107.6	107.4	189.6 - 239.4	213.2	212.8
16	113.6 - 128.2	120.5	120.4	239.4 - 301	269	268.9
17	128.2 - 147.6	137.4	137.3	301 - 380.1	338.7	337.5
18	147.6 - 176.2	160.9	160.3	380.1 - 492.8	433.3	431.7
19	176.2 - 233.4	200.4	198.4	492.9 - 662.9	573.9	571.8
20	233.4 - 7,391	414.4	307.1	662.9 - 957.4	793.9	783.9
21				957.6 - 1,628	1,228.6	1,191.4
22				1,628 - 117,649	4,174.6	2,590

Table 3: Effects of Winsorizing by Income Bin

Income Bin	Mean Percent of Income Given	10% Winsor		20% Winsor		30% Winsor	
		Mean	% Affected	Mean	% Affected	Mean	% Affected
1	33.60	1.30	7.91	2.01	5.69	2.56	4.33
2	1.82	1.12	4.78	1.43	2.26	1.58	0.87
3	2.07	1.46	5.69	1.76	1.58	1.87	0.75
4	1.75	1.32	4.40	1.54	1.05	1.62	0.57
5	1.94	1.33	4.29	1.54	0.94	1.60	0.38
6	1.79	1.34	4.22	1.57	1.01	1.65	0.64
7	1.65	1.33	3.84	1.50	0.90	1.56	0.49
8	1.76	1.50	4.46	1.67	0.57	1.70	0.23
9	1.72	1.50	4.17	1.67	0.49	1.69	0.26
10	1.79	1.55	3.99	1.73	0.68	1.76	0.23
11	1.76	1.59	4.17	1.73	0.34	1.75	0.11
12	1.82	1.58	3.69	1.69	0.34	1.71	0.19
13	1.72	1.59	3.01	1.70	0.38	1.72	0.04
14	1.77	1.63	3.35	1.75	0.30	1.76	0.08
15	1.82	1.66	3.65	1.78	0.30	1.81	0.19
16	1.99	1.73	3.47	1.86	0.60	1.90	0.41
17	1.83	1.74	2.82	1.82	0.19	1.83	0.08
18	1.96	1.77	3.28	1.89	0.45	1.93	0.30
19	2.04	1.78	2.82	1.89	0.45	1.92	0.19
20	1.71	1.60	2.18	1.69	0.34	1.71	0.04

Table 4: Higher-Income Giving
(IRS Statistics of Income)

IRS SOI Income Groups	Pr (Giving>0)	Mean Amount Given	Mean Percent of Income Given
$200,000 - $500,000	85.2%	$6,302	1.97%
$500,000 - $1,000,000	88.5%	$16,831	2.32%
$1,000,000 - $1,500,000	86.4%	$36,088	2.78%
$1,500,000 - $2,000,000	86.8%	$53,300	2.91%
$2,000,000 - $5,000,000	88.3%	$108,119	3.44%
$5,000,000 - $10,000,000	92.2%	$308,269	4.32%
$10,000,000+	94.6%	$2,661,195	8.57%

Notes: IRS Statistics of Income (SOI) tables report the number of filers in each income group as well as the number of itemizers. Means in this table are calculated assuming that non-itemizers do not make any donations.

Appendix

Figure A.1: OLS, Unconditional Winsorized Percent of Giving, by Income

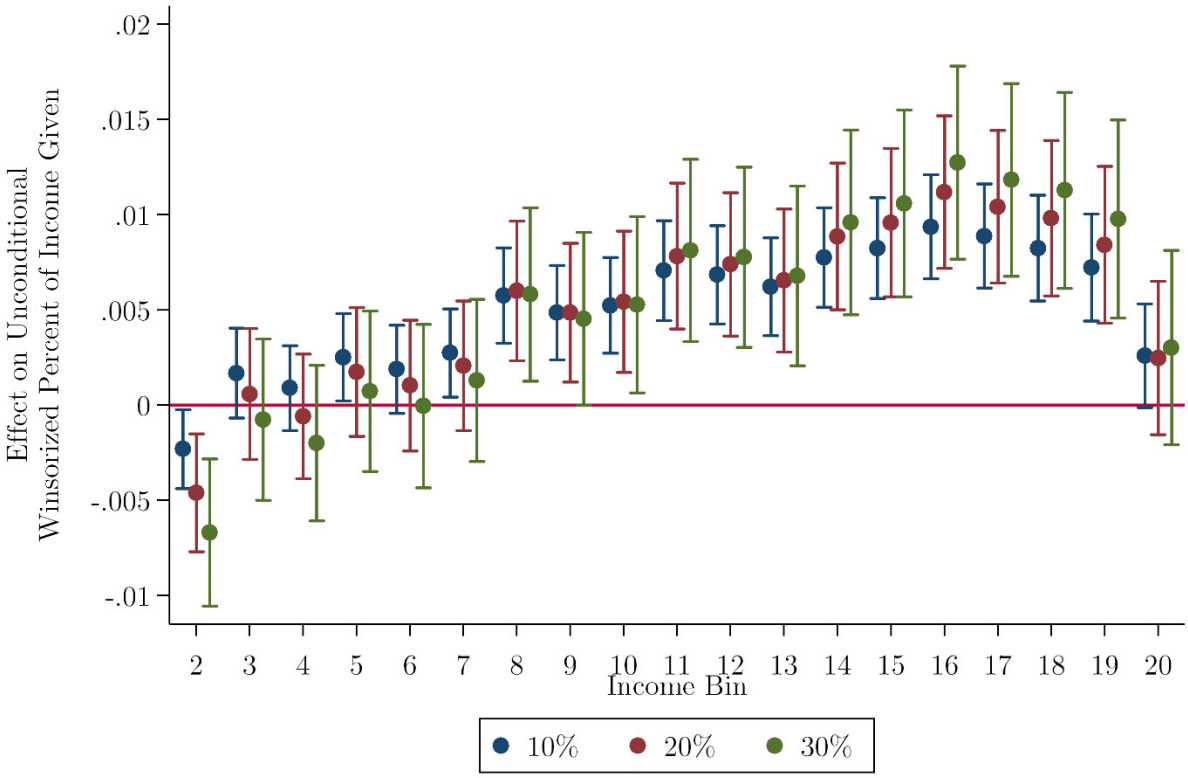

Notes: This figure displays the coefficients for each income bin from OLS regressions for three levels of winsorizing: 10%, 20%, 30%. Estimates combine the extensive and intensive margins. Winsorizing caps the maximum possible percent of income given for a household at 10%, 20%, or 30%. Households giving a larger proportion of their income are assigned a value equal to the maximum. All specifications include income bin and year indicators; the lowest income bin is omitted for comparison. Household characteristics described in Section 3.2 are included. Error bars represent the 95% confidence interval. Robust standard errors are clustered at the household level.

Figure A.2: OLS, Unconditional Winsorized Percent of Giving, by Wealth

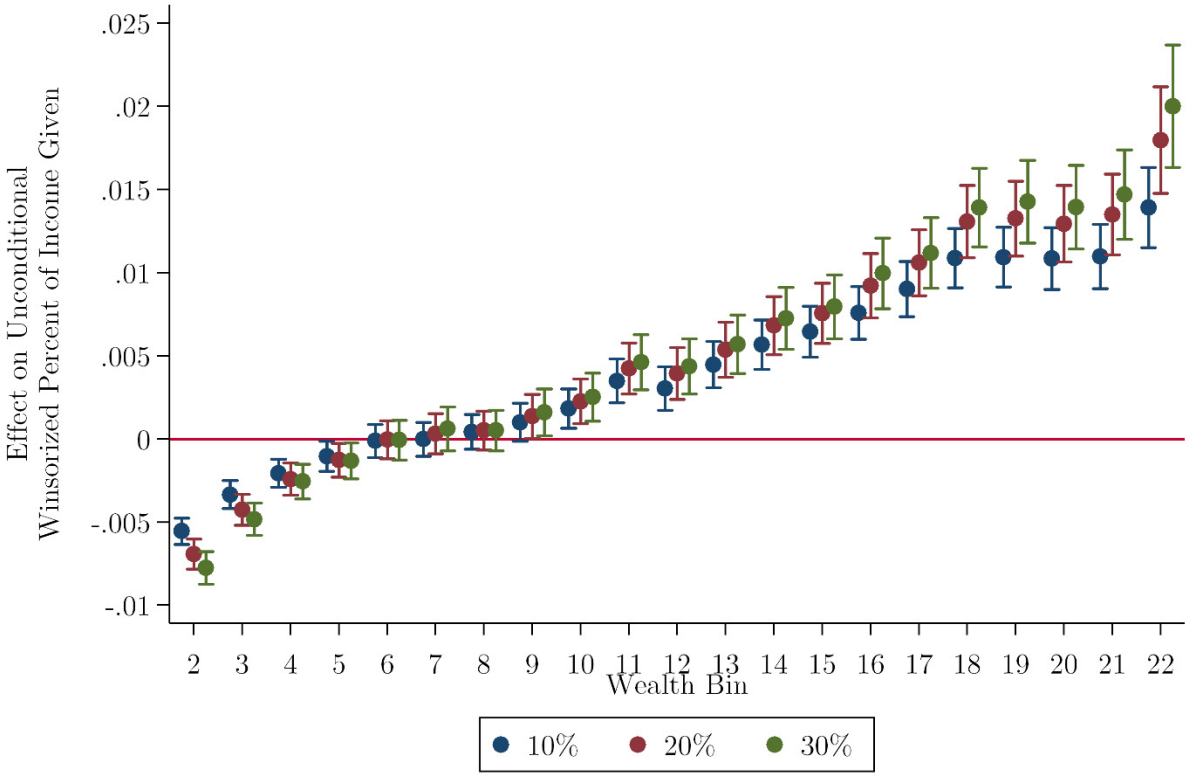

Notes: This figure displays the coefficients for each wealth bin from OLS regressions for three levels of winsorizing: 10%, 20%, 30%. Estimates combine the extensive and intensive margins. Winsorizing caps the maximum possible percent of income given for a household at 10%, 20%, or 30%. Households giving a larger proportion of their income are assigned a value equal to the maximum. All specifications include wealth bin and year indicators; the lowest wealth bin (negative wealth) is omitted for comparison. Household characteristics described in Section 3.2 are included. Error bars represent the 95% confidence interval. Robust standard errors are clustered at the household level.

Figure A.3: Fixed Effects, Unconditional Winsorized Percent of Giving, by Income

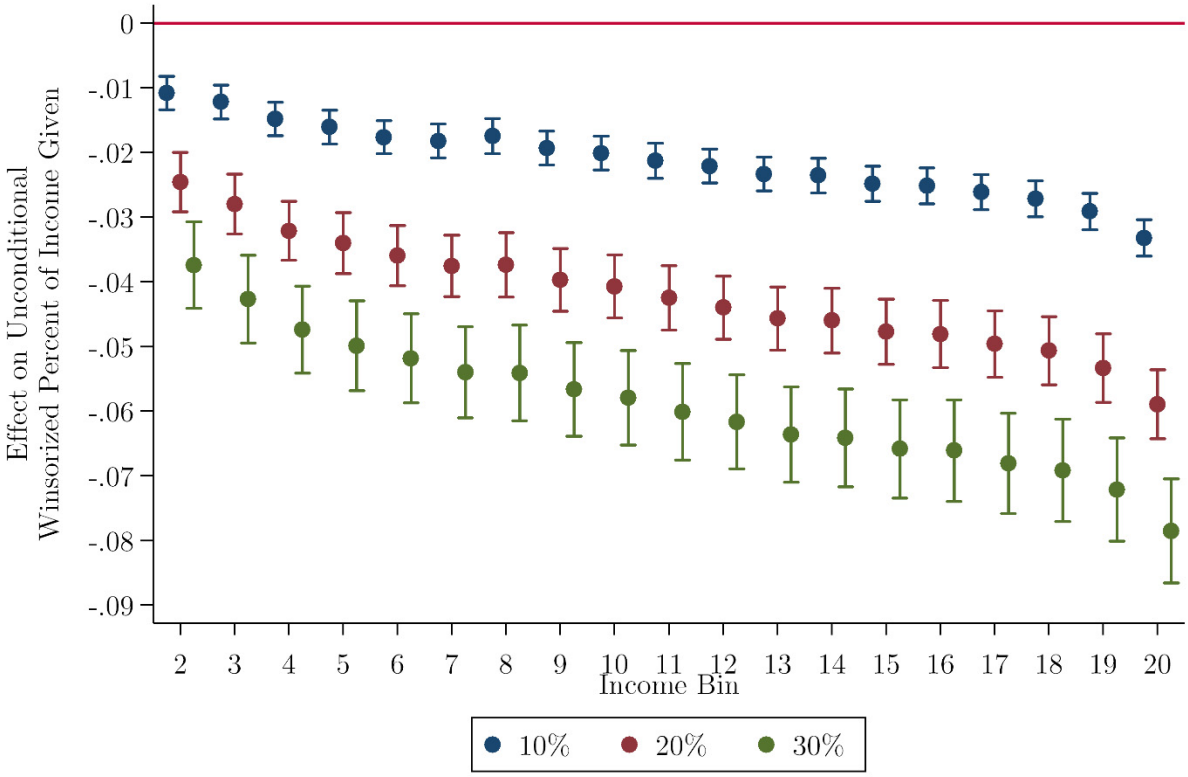

Notes: This figure displays the coefficients for each income bin from fixed effects regressions for three levels of winsorizing: 10%, 20%, 30%. Estimates combine the extensive and intensive margins. Winsorizing caps the maximum possible percent of income given for a household at either 10%, 20%, or 30%. Households giving a larger proportion of their income are assigned a value equal to the maximum. All specifications include income bin and year indicators; the lowest income bin is omitted for comparison. Time-varying household characteristics described in Section 3.2 are included. Error bars represent the 95% confidence interval. Robust standard errors are clustered at the household level.

Figure A.4: Fixed Effects, Unconditional Winsorized Percent of Giving, by Wealth

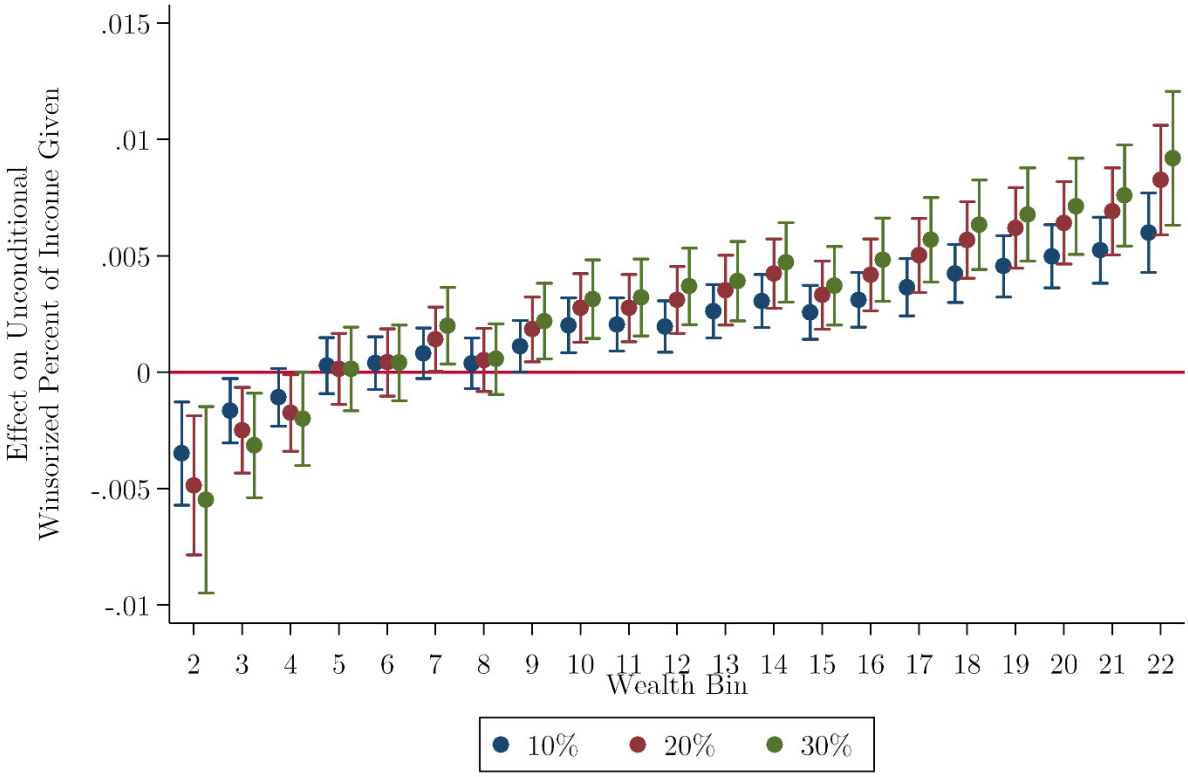

Notes: This figure displays the coefficients for each wealth bin from fixed effects regressions for three levels of winsorizing: 10%, 20%, 30%. Estimates combine the extensive and intensive margins. Winsorizing caps the maximum possible percent of income given for a household at either 10%, 20%, or 30%. Households giving a larger proportion of their income are assigned a value equal to the maximum. All specifications include wealth bin and year indicators; the lowest wealth bin (negative wealth) is omitted for comparison. Time-varying household characteristics described in Section 3.2 are included. Error bars represent the 95% confidence interval. Robust standard errors are clustered at the household level.