

NBER WORKING PAPER SERIES

THE IMPORTANCE OF BEING MARGINAL:
GENDER DIFFERENCES IN GENEROSITY

Stefano DellaVigna
John A. List
Ulrike Malmendier
Gautam Rao

Working Paper 18748
<http://www.nber.org/papers/w18748>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2013

This paper will appear as part of the American Economic Review Papers and Proceedings paper, May 2013. Many thanks to Alec Brandon, Sheng Li, David Novgorodsky, and Vera te Velde for able research assistance. We are grateful to Rachel Croson and Yan Chen as well as to seminar participants at UC Berkeley, CalTech, Chicago, and the ASSA 2013 and ESA 2012 meetings for valuable comments. Funding for this research was obtained from funds of the Becker Friedman Institute at the University of Chicago. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2013 by Stefano DellaVigna, John A. List, Ulrike Malmendier, and Gautam Rao. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Importance of Being Marginal: Gender Differences in Generosity
Stefano DellaVigna, John A. List, Ulrike Malmendier, and Gautam Rao
NBER Working Paper No. 18748
February 2013
JEL No. C93,D64,H4

ABSTRACT

Do men and women have different social preferences? Previous findings are contradictory. We provide a potential explanation using evidence from a field experiment. In a door-to-door solicitation, men and women are equally generous, but women become less generous when it becomes easy to avoid the solicitor. Our structural estimates of the social preference parameters suggest an explanation: women are more likely to be on the margin of giving, partly because of a less dispersed distribution of altruism. We find similar results for the willingness to complete an unpaid survey: women are more likely to be on the margin of participation.

Stefano DellaVigna
University of California, Berkeley
Department of Economics
549 Evans Hall #3880
Berkeley, CA 94720-3880
and NBER
sdellavi@econ.berkeley.edu

Ulrike Malmendier
Department of Economics
549 Evans Hall # 3880
University of California, Berkeley
Berkeley, CA 94720-3880
and NBER
ulrike@econ.berkeley.edu

John A. List
Department of Economics
University of Chicago
1126 East 59th
Chicago, IL 60637
and NBER
jlist@uchicago.edu

Gautam Rao
Department of Economics
549 Evans Hall # 3880
University of California, Berkeley
Berkeley, CA 94720-3880
grao@berkeley.edu

In modern economies, private organizations are relied on to provide important public goods and services. Indeed, since 1971 individual contributions to charitable causes have increased from roughly \$130 billion to nearly \$300 billion, or about 2% of GDP. Despite much interest in understanding the ‘whys?’ and ‘whats?’ of giving, however, fundraisers and academics alike have faced some persistent puzzles when trying to predict which individual will give for which cause. The predictive power of demographic and other individual characteristics varies widely across data sets and studies. In the case of gender, the focus of our study, there are conflicting views on which is the more charitable gender. In the laboratory, using variants of the dictator game, Bolton and Katok (1995) find no evidence of gender differences in generosity, while Eckel and Grossman (1998) find that women share twice as much on average.

The same puzzle exists in field data. Looking at the statistics of blood donation, women sometimes are overrepresented (e.g., among first-time Norwegian blood donors in Misje, Bosnes, and Heier, 2010) and sometimes underrepresented (e.g., in most European countries in the meta analysis of Bani and Guissani, 2010) relative to men.

We provide new evidence from a field experiment that can help reconcile prior conflicting findings. We argue that seeming inconsistencies in gender-specific patterns of pro-social behavior reflect, at least partly, distributional differences in altruism between genders. Even in cases where men and women do not differ in their mean inclination to give, differences in heterogeneity can lead to a higher share of women being on the margin of giving. Hence, small differences in the giving request (such as in the cost of saying no) can have sizeable effects on women’s pro-social behavior relative to men. Hence, it is important to estimate the within-gender heterogeneity in social preferences to better understand the gender differences.

Our analysis builds on recent attempts to tie models of pro-social behavior more closely to

empirical tests, obtaining structural estimates of the underlying preferences. Specifically, DellaVigna et al. (2012) use a door-to-door fund-raising campaign and survey to disentangle the importance of warm-glow altruism versus social pressure in charitable giving. Their estimates suggest that the social pressure cost of saying no to a solicitor plays an important role in high-pressure giving requests such as door-to-door campaigns. In this paper, we exploit the rich design in DellaVigna et al. (2012) to estimate the distribution of social preferences by gender.

DellaVigna et al. (2012) compare the incidence of giving in a control treatment, in which subjects receive an unannounced door-to-door visit, to two flyer treatments, in which subjects are notified a day in advance of the upcoming door-to-door campaign via a flyer on the door knob. The flyer treatments allow donors who give due to altruistic motives to sort in; and they allow donors who give due to social pressure to sort out. In one of the two flyer treatments, the flyer also includes an opt-out box, which makes it easy to avoid the solicitor (the solicitor does not contact the household if the box is checked). The main findings are that (i) the simple flyer lowers the share answering the door, relative to the control group, but does not affect the share giving; and that (ii) the opt-out option lowers both the share answering the door and the share of individuals giving. These findings suggest that both altruism and social pressure are at play. The desire to avoid social pressure explains the drop in giving in the opt-out treatment.

In this paper, we decompose these findings by gender of the respondents. In Figure 1 we plot for each treatment the share of households in which a male person answers and gives to the charity out of all the households contacted; similarly, we compute a share of females answering and giving to the charity. The sum of the two shares equals approximately the unconditional share of givers out of all the individuals contacted.² As Figure 1 shows, the shares of male and

² For a small fraction of respondents, gender was not recorded. Notice also that we do not observe the gender of the potential giver in the case of households not answering the door.

female givers are about the same in the baseline treatment. Given that the shares of males and females answering the door (not shown) are also similar in this treatment, we conclude that generosity in the baseline treatment does not differ by gender. The shares remain similar in the simple flyer treatment. In the opt-out treatment, however, women are significantly *less* likely to give. The decrease in male giving in the opt-out treatment is instead small.




Figure 1. Share of households contacted giving to a charity, by gender and treatment

These results provide yet another example of the seemingly inconsistent findings in giving behavior: In a setting where individuals cannot (easily) avoid a giving request, we find no gender differences; in a setting where individuals can sort out at a low cost, instead, the data suggests significant gender differences in giving behavior. Considering these findings jointly, however, we conclude that women are more responsive to a simple avoidance strategy, the opt-out box.

We estimate the social preference parameters separately by gender. Using all the empirical moments in the data (not just those displayed in Figure 1), and allowing the key parameters to differ by gender, we find a significant difference not only in the mean, but also in the variance of

the gender-specific altruism distributions. The implied distributions indicate that a higher share of women is on the margin of giving, and hence responsive to a low-cost opportunity to opt out (since marginal givers prefer not to be asked in the first place). We also find evidence of a similar pattern with respect to the pro-social willingness to answer a survey request.

Our findings point to the importance of considering differences in the overall distribution of social preferences and, in particular, marginality. When put under pressure, women may give more, and contribute more to public goods because they are more likely to be on the margin, and hence sensitive to an extra push. But they may say no if given a simple option to do so.

The finding in this paper are consistent with the hypothesis that women are more malleable or more sensitive to social cues in determining appropriate behavior (see, e.g., Kahn et al., 1971; Croson and Gneezy, 2009).³ Our findings are not inconsistent with the findings cited above that suggest that women are more generous than men. A natural interpretation is that the presence of a larger share of marginal givers leads women to give more in certain situations, but not in others. We would like to emphasize, though, that the results in this paper should be seen as suggestive and that more evidence will be necessary. Finally, this paper highlights the benefits of the literature on Structural Behavioral Economics. While the reduced-form findings on gender differences in giving are of first order importance, it is the estimation of the underlying giving preferences which suggests a possible explanation of differences in giving across settings.

I. Framework and Experimental Design

³ Our findings also complement the price sensitivity of giving established by Andreoni and Vesterlund (2001), who show that women are more generous than men when it is relatively expensive to give, but that men begin to give more than women as the price of giving decreases.

Framework. Consider a two-stage game between a solicitor and a potential giver (solicitee).⁴ In the first stage, the solicitee may receive a warning of the upcoming solicitation via a flyer at the door, which she notices with probability $r \in (0,1]$. In the second stage, the solicitor visits the home. The solicitee opens the door with probability $h \in [0,1]$. If she did not notice the flyer (or did not receive one), then h is equal to $h_0 \in (0,1)$. If she noticed the flyer, then she can adjust the probability to h at a cost $c(h)$, with $c(h_0) = 0$, $c'(h_0) = 0$, and $c''(\cdot) > 0$.

If the solicitee opens the door, then she donates $g \geq 0$. If she does not open the door, there is no donation ($g = 0$). In our set-up, a solicitee of gender $i \in \{female, male\}$ has utility given by:

$$U^i(g) = u^i(W - g) + a^i v^i(g, G) - s^i(g). \quad (1)$$

The overall utility U^i of an individual of gender i is composed of three terms. (For simplicity, we suppress the index for each individual.) First, private consumption is denoted by u^i , which includes wealth W minus the individual's donation g . We model this private utility as satisfying standard properties: $(u^i)' > 0$ and $(u^i)'' \leq 0$. The second term in (1), $a^i v^i$, allows for pure and impure altruism (warm glow). In the case of pure altruism, the agent cares about the total contributions to the charity, $G + g$, where G is the giving of others. In the case of impure altruism, the agent cares about the warm glow from giving g , and v^i does not depend on the giving of others. The altruism parameter $a^i \geq 0$ captures the intensity of the warm glow (the case $a^i < 0$ captures the possibility of spite). Importantly, a^i is assumed to be heterogeneous across people of gender i , with a distribution F^i .

⁴ We only summarize the framework here, directing the interested reader to the details and the slightly more general model in DellaVigna et al. (2012).

The third component of (1) is social pressure. The solicitee absorbs a utility cost $s^i(g) = S^i \cdot (g^n - g) \cdot 1_{g < g^n} \geq 0$ for not giving (or for giving a small amount), with this cost decreasing linearly in g . The giver does not incur a social pressure cost if she is away from home during the visit. The special case of $S^i = 0$ (no social pressure) and $a^i = 0$ (no altruism or warm glow) represents the standard model.

The model yields several testable implications. When altruism dominates social pressure, the flyer increases home presence and giving relative to the control group since some agents seek to meet the solicitor. When social pressure dominates, instead, the flyer treatments, and especially the opt-out ones, lead to lower rates of answering the door and of donations. We allow these effects to differ by gender.

We also model the decision to respond to a survey request, which is a form of giving, namely of giving time for a survey. We assume that individuals of gender i receive a utility sv^i (which could be positive or negative) from completing a 10-minute survey for no monetary payment. In addition, individuals receive utility from a payment m^i for completing the survey, and receive disutility from the time cost t^i of the survey, both of which are deterministic. The overall utility from completing a survey hence is $sv + m - t$. We assume that each individual has a pro-social value sv^i drawn from a normal distribution.

We structurally estimate the model to provide evidence on the social preference parameters for men and women.

Experimental Design. Our design combines two elements. First, we raise money in a door-to-door fundraising drive for two charities: La Rabida Children's Hospital and the East Carolina Hazard Center (ECU). Second, we conduct surveys of varying lengths (and varying

monetary incentives as inducements) to estimate the elasticity of the presence at home and of the response rate to incentives.

In the control treatments, solicitors knock on the door or ring the bell and, if they reach a person, proceed through the script. In case of the charity solicitation, they inform the household about the charity (La Rabida or ECU) and ask if they are willing to make a donation. In case of the survey solicitation, they inquire whether the household member is willing to answer a survey about charitable giving. The solicitor informs the household member about the duration of the survey (5 or 10 minutes) and about the payment for completing the survey, if any (\$10, \$5, or none). In the flyer treatments, the solicitor's script is identical, but in addition a different solicitor leaves a flyer on the doorknob the day before the solicitation. The professionally prepared flyer indicates the time of the upcoming fund-raising (or survey) visit within a one-hour time interval. In the treatments with opt-out, these flyers include a box that says: "Check this box if you do not want to be disturbed." If the solicitors find the box checked, they do not knock on the door.


The field experiment took place on Saturdays and Sundays between April 2008 and October 2008. We employed 92 solicitors and surveyors, mostly undergraduate students at the University of Chicago.

II. Structural Estimates

We estimate the parameters of the model using a minimum distance estimator following the same procedure as in DellaVigna et al. (2012), to which we refer for details. We allow some of the key parameters to differ by gender, namely the parameters determining the distribution of altruism a towards the charities: the mean and the standard deviation of altruism. We also allow for a different social pressure cost of turning down a solicitor for males and females. Finally, we

model the willingness to complete an unpaid 10-minute survey as a normal distribution, with mean and standard deviation that differ by gender.

The resulting estimates are reported in Table 1 of the Online Appendix. Among the most relevant parameters, the standard deviation of the altruism distribution is smaller for women (17.4) than for men (19.5); in addition, there is a sizeable difference in the mean altruism towards the out-of-state charity (ECU). Figure 2 plots the implied density of the altruism distribution for the two charities, separately for each gender. Noted is the threshold for positive utility from giving.⁵ As Figure 2 shows, for the out-of-state charity (ECU) the share of types that are on the margin of giving – that is the types at the threshold – is significantly larger for women. There is a smaller difference for the in-state charity. The larger density of estimated types at the margin for women implies that women will be more responsive to shifts in features of the environment, or in the cost of sorting out.



⁵ The actual threshold for giving is $a = (1-S)G$ and thus depends on the estimated social pressure cost which differs by charity and by gender; this threshold lies between 0 (plotted in the figure) and $G (= 10)$, the threshold with no social pressure.

Figure 2. Implied distribution of altruism towards the two charities, by gender

While the emphasis so far is on giving of money, what about giving of time? We use the field experiment on door-to-door survey completion to estimate the willingness to complete an unpaid 5-minute survey. Figure 1 in the Online Appendix shows that the share of women completing the survey decreases significantly from the flyer treatment to the flyer with opt-out, consistent with the charitable giving results. Table 1 in the Online Appendix shows that the estimated standard deviation of the pro-social utility of completing a survey is again smaller for females (26.3) than for males (34.0), although the difference is not statistically significant at conventional levels. Figure 3 displays the implied distribution of pro-social utility from survey completion, which again shows that women are more likely to be at the marginal point.⁶

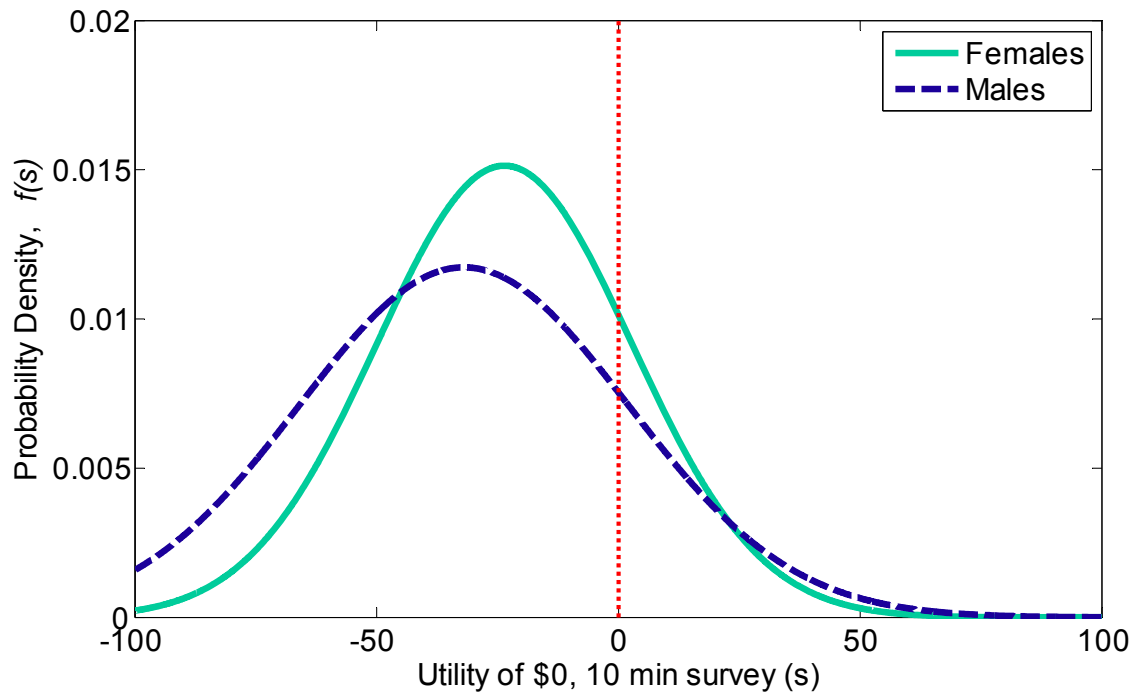


Figure 3. Implied distribution of willingness to complete a 5-minute unpaid survey.

⁶ As in the case of giving money, the figure plots the threshold for positive willingness to do a survey; the threshold for survey completion is $-S$.

III. Conclusions

This study uncovers an important relationship between gender and giving patterns: there are gender differences in social preferences, but it is important to go beyond considering differences in means –important gender differences may be at the margin. This leads women to give more in certain situations, but not in others, and also to be more sensitive to social cues.

Our study revolves around an experimental design that is tightly linked to a theory of altruism and social pressure. The results naturally permit to improve our understanding of the quantitative importance of each determinant of charitable giving. Differentiating by gender reveals a novel explanation for seemingly contradictory findings in previous literature, and our methodology is applicable to other determinants of giving.

References

- Andreoni, Jim and Lise Vesterlund. 2001. "Which is the fair sex? Gender differences in altruism." *Quarterly Journal of Economics* 116(4): 293-312.
- Bani, Marco and Barbara Guissani. 2010. "Gender differences in giving blood: a review of the literature." *Blood Transfusion* 8(4): 278-287.
- Bolton, Gary and Elena Katok. 1995. "An experimental test for gender differences in beneficent behavior." *Economics Letters* 48(1995): 287-292.
- DellaVigna, Stefano, John List, and Ulrike Malmendier. 2012. "Testing for Altruism and Social Pressure in Charitable Giving", *Quarterly Journal of Economics*, 127, 1-56.
- Eckel, Catherine C. and Philip J. Grossman. 1998. "Are Women Less Selfish Than Men?: Evidence From Dictator Experiments." *Economic Journal* 108(448): 726-735.
- Misje, A. H., V. Bosnes, and H. E. Heier. 2010. "Gender differences in presentation rates, deferrals and return behaviour among Norwegian blood donors." *Vox Sanguinis* 98: e241-e248.

Online Appendix for “The Importance of Being Marginal: Gender Differences in Generosity”

Stefano DellaVigna, John List, Ulrike Malmendier, Gautam Rao

January 14, 2013

This appendix describes the structural estimation procedure underlying the quantitative estimates of gender-specific social preferences presented in the main paper. The estimation procedure generalizes DellaVigna et al. (2012) to allow for gender-specific parameters. In the theoretical model, we suppress gender indicators for notational simplicity; we then spell out in the description of the estimation procedure which estimates are gender-specific. We also report the reduced form results by gender of completion of an unpaid survey.

Completion of an unpaid survey. Figure 1 reports the reports the share completing an unpaid 5-minute survey in 2009 for the three treatments – baseline, flyer, and flyer with opt-out. The bars for each gender report the share of all households contacted which answer the door, agree to complete the survey and are of a specific gender. Hence, the sum of the male and female bars equals (up to a small share of respondents with unreported gender) the unconditional share completing the survey. The patterns in Figure 1 indicate that in the control group women are less likely to complete the survey than male, but are somewhat more likely to complete the survey in the flyer treatment, indicating a higher share of women sorting in. Finally, in the opt-out treatment the share of

women completing the survey is lower than the share of men, consistent with the results of the charitable giving fund-raising. This latter result indicates higher sorting out by women when the cost of doing so is small (checking a box).

Theoretical Framework

We consider a two-stage game between a potential giver and a solicitor. In the first stage, the giver may receive a flyer of the upcoming visit and, if so, notices the flyer with probability $r \in (0, 1]$. In the second stage, the solicitor visits the home. The giver opens the door with probability h . If she did not notice the flyer (or did not receive one), h is equal to a baseline probability $h_0 \in (0, 1)$. If she noticed the flyer, she can adjust the probability to $h \in [0, 1]$ at a cost $c(h)$, with $c(h_0) = 0$, $c'(h_0) = 0$, and $c''(\cdot) > 0$. That is, the marginal cost of small adjustments is small, but larger adjustments have an increasingly large cost.

Charity donation solicitation. If the giver is present, she donates an amount $g \geq 0$. If she is absent, there is no donation ($g = 0$). A giver j of gender $i \in \{male, female\}$ has utility

$$U^i(g_j) = u(W - g_j) + a_j^i v^i(g_j, G^i) - s^i(g_j). \quad (1)$$

In the discussion that follows, we suppress the indices for individual as well as gender. The utility of private consumption, u , is derived from the pre-giving wealth W minus the donations given to the solicitor (g). The private utility satisfies standard properties: $u'(\cdot) > 0$ and $u''(\cdot) \leq 0$.

The utility from giving to the charity is $av(\Gamma + g)$, where the parameter Γ governs the concavity of the altruism function. In the case of pure altruism, $\Gamma \equiv G$ is the amount given by others to the charity. Then, we can think of $v(G + g)$ as the production function of the charity, which is increasing in the donation g but has decreasing returns: $v'_g(\cdot, \cdot) > 0$, $v''_{g,g}(\cdot, \cdot) < 0$, and $\lim_{g \rightarrow \infty} v'(g, \cdot) = 0$.

The parameter $a \geq 0$ denotes the level of altruism, and the overall utility from giving is $av(G + g)$. More generally, in the case of impure altruism (warm glow), Γ need not equal G .

The third element in the utility function is social pressure. The giver of gender i pays a utility cost $s(g) = S \cdot (g^s - g) \cdot 1_{g < g^s} \geq 0$ for not giving or only a giving small amount to the solicitor. The cost is highest for the case of no donation ($s(0) = Sg^s$), decreases linearly in g , and is zero for donations of g^s or higher. The giver does not incur a social pressure cost if she is away from home during the visit. The special case of $S = 0$ (no social pressure) and $a = 0$ (no altruism or warm glow) represents the standard model. We further assume that the giver is aware of her own preferences and rationally anticipates her response to social pressure.

Survey solicitation. Individuals receive a utility sv (which could be positive or negative) from completing a 10-minute survey for no monetary payment. In addition, individuals receive utility from a payment m for completing the survey, and receive disutility from the time cost t of the survey, both of which are deterministic. Assuming (locally) linear utility, we add these terms to obtain the overall utility from completing a survey: $sv + m - t$. We denote by S^{sv} the social pressure cost of saying no to a survey request. The agent undertakes the survey if $sv + m - t$ is larger than $-S^{sv}$. The threshold $\bar{sv}^{m,t} = -S^{sv} - (m - t)$ is the lowest level of sv such that individuals will agree to complete the survey if asked. An increase in the social pressure S^{sv} or in the pay m , or a decrease in the cost of time t will lower the threshold and hence increase the probability of survey completion. The decision problem of staying at home conditional on receiving a notice is

$$\max_{h \in [0,1]} h \max(sv + m - t, -S^{sv}) - \frac{(h - h_0)^2}{2\eta}. \quad (2)$$

Taking into account corner solutions for h^* , this leads to a solution for the probability of being at home: $h^* = \max[\min[h_0 + \eta \max(sv + m - t, -S^{sv}), 1], 0]$.

Assumptions. To estimate the model, we impose the following additional assumptions.

1. We assume that the homes approached are of either male or female “type”, with a share p of female homes. Thus, we abstract away from the presence of multiple household members who might engage in collective decision making about who should answer the solicitor’s knock. The parameter p is identified as the share of females answering the door when the solicitor’s visit is unanticipated. This same share p is assumed to apply to homes in all the treatments - in the opt-out flyer treatment, for example, we assume that a share p of the homes to which the flyer is delivered are of the female type.
2. The private utility of consumption is linear, $u(W - g) = W - g$. This assumption is justified by the local linearity implied by a model of expected utility.
3. The parameter for altruism towards the charities, a , comes from a normal distribution with mean μ and variance σ^2 . In the estimation, the distributions are allowed to differ across genders by mean (reflecting the different nature of the two charities) and variance; however, we impose that the variance of altruism be the same for the two charities.¹
4. The utility sv of completing a 10-minute survey is assumed to be normally distributed with parameters μ^{sv} and σ^{sv} . We allow sv to be negative for households that dislike doing surveys without compensation. Both the mean and the variance are allowed to differ across genders.

¹Results of estimations allowing the variance of altruism to also differ across charities are available upon request. The results for La Rabida are similar to those reported below, but the variance for ECU is imprecisely estimated.

5. The altruism function for charitable giving is $av(g, G) = a \log(\Gamma + g)$, where the parameter Γ governs the concavity of the altruism function for $a > 0$: a large Γ implies that the marginal utility of giving, given by $a/(\Gamma + g)$, declines only slowly in the individual giving g , consistent with pure altruism—the individual cares about the overall donation and her individual giving is only a small part. A small Γ instead indicates that the marginal utility diminishes steeply with the individual giving, more consistent with warm glow. For the results presented here, we fix the value of $\Gamma = 10$, close to the estimated value in DellaVigna et al (2012), and assumed identical across men and women.
6. The social pressure cost S incurred from saying no to the solicitor is allowed to differ across genders and charities, but is assumed to be homogeneous within genders.
7. The level of giving g^S from which on there is no social pressure cost is \$10 (the median donation), for both men and women.
8. The cost of leaving home $c(h)$ is symmetric around h_0 and quadratic: $c(h) = (h - h_0)^2 / 2\eta$. For the estimates presented in the paper, the elasticity η is assumed to be equal across genders.²

Estimated Parameters. The vector of parameters ξ that we estimate are: (i) p , the share of female homes; (ii) h_0^{2008} and h_0^{2009} —the probabilities of opening the door in the 2008 and 2009 no-flyer treatments; (iii) r —the probability of observing (and remembering) the flyer, assumed equal across genders; (iv) η —the responsiveness of the probability of opening the door to the desirability of being at home, assumed equal across genders; (v) μ_a^{ch} and σ_a (where $ch = LaR, Ecu$)—the mean and standard deviation of the normal distribution F from which the

²Preliminary estimations available upon request suggest that allowing η to vary across gender does not change the main conclusions.

altruism parameter a is drawn; we allow the mean μ_a^{ch} to differ across genders; (vi) Γ —the curvature of the altruism function, which is assumed to be the same for the two charities and for men and women; (vii) S^{ch} ($ch = LaR, Ecu$)—the social pressure cost associated with a donation request, allowed to differ by gender; the table displays the social pressure cost associated with giving zero, $Sg^S = 10S$; (viii) μ^{sv} and σ^{sv} —the mean and standard deviation of the utility of completing an unpaid 10-minute survey, which differ across genders; (ix) v^{sv} —the value of one hour of time spent completing a survey, assumed equal across genders; and (x) S^{sv} —the social pressure associated with saying no to the survey request, differing across men and women.

Estimation Method. To estimate the model, we use a minimum-distance estimator. Denote by $m(\xi)$ the vector of moments predicted by the theory as a function of the parameters ξ , and by \hat{m} the vector of observed moments. The minimum-distance estimator chooses the parameters $\hat{\xi}$ that minimize the distance $(m(\xi) - \hat{m})' W (m(\xi) - \hat{m})$, where W is a weighting matrix. As a weighting matrix, we use the diagonal of the inverse of the variance-covariance matrix. Hence, the estimator minimizes the sum of squared distances, weighted by the inverse variance of each moment.³ To calculate the theoretical moments for the probability of opening the door and the probability of giving, we use a numerical integration algorithm based on adaptive Simpson quadrature, implemented in Matlab as the *quad* routine.

Moments. As moments $m(\xi)$ we use the probabilities of taking the various actions (answering the door, giving, completing the survey), each broken down by gender. Note that we do not observe the gender for households who do not answer the door, or who check the opt-out box. Therefore, as empirical moments we use the share that has a certain gender, out of the whole population contacted

³Given the large number of moments, weighting the estimates by the inverse of the full variance-covariance matrix is computationally difficult.

in that treatment. For example, the probability of a female answering the door in treatment i is the number of households in which a female answered the door out of all households contacted in treatment i .⁴

The moments are as follows (where $j = F, NF, OO$ and $ch = LaR, Ecu$): (i) the probability of opening the door in the various charity treatments ($P(H)_j^{ch}$); (ii) the probability of checking the opt-out box in the Opt-Out treatment ($P(OO)_{OO}^{ch}$); (iii) the unconditional probability of giving in the various charity treatments ($P(G)_j^{ch}$); (iv) the probability of giving an amount of money in different ranges ($P(0 < G < 10)_j^{ch}$, $P(G = g^s = 10)_j^{ch}$, $P(10 < G \leq 20)_j^{ch}$, $P(20 < G \leq 50)_j^{ch}$, and $P(G > 50)_j^{ch}$); (v) the probability of opening the door in the various survey treatments k (with varying dollar amounts, minutes, and flyer conditional), $P(H)_k^{sv}$, run in 2008 and in 2009; (vi) the unconditional probability of completing the survey in the various survey treatments, $P(SV)_k^{sv}$, run in 2008 and in 2009; and (vii) the probability of checking the opt-out box in the survey Opt-Out treatments ($P(OO)_k^{sv}$). The corresponding empirical moments \hat{m} are estimated in a first stage model using the same controls as in the reduced form regressions in DellaVigna et al. (2012), including solicitor fixed effects and day-time fixed effects.

To calculate the method of minimum distance estimate, we employ a common sequential quadratic programming algorithm (Powell, 1983) implemented in Matlab as the *fmincon* routine. We impose the following constraints: $S^{ch}, S^{sv} \geq 0$ (social pressure non-negative), $\sigma^{ch}, \sigma^{sv} > 0$ (positive standard deviation of altruism), $h_0^{2008}, h_0^{2009}, r \in [0, 1]$ (probabilities between zero and one), and $\eta \in [0, 9999]$ (finite elasticity of home presence). We begin each run of the optimization routine by randomly choosing a starting point, drawn from a uniform

⁴Since we do not observe the share opting out disaggregated by gender, the moment $P(OO)$ is not split by gender. Theoretically, it is calculated as the sum of the shares of females and males choosing to opt out, weighted by the estimated shares p and $1 - p$ of female and male households.

distribution over the permitted parameter space. The algorithm determines successive search directions by solving a quadratic programming sub-problem based on an approximation of the Lagrangian of the optimization problem. To avoid selecting a local minima, we choose the run with the minimum squared distance of 500 runs.⁵

Under standard conditions, the minimum-distance estimator using weighting matrix W achieves asymptotic normality, with estimated variance

$$\text{Var}=(\hat{G}'W\hat{G})^{-1}(\hat{G}'W\hat{\Lambda}W\hat{G})(\hat{G}'W\hat{G})^{-1}/N$$

where $\hat{G} \equiv N^{-1} \sum_{i=1}^N \nabla_{\xi} m_i(\hat{\xi})$ and $\hat{\Lambda} \equiv \text{Var}[m(\hat{\xi})]$ (Wooldridge, 2002). We calculate $\nabla_{\xi} m(\hat{\xi})$ numerically in Matlab using an adaptive finite difference algorithm.

Identification. While the parameters are estimated jointly, it is possible to describe the main sources of identification of individual parameters. The share of female households p together with the baseline probabilities of answering the door, h_0^{2008} and h_0^{2009} , are identified by the observed probabilities of opening the door in treatments without flyer. The probability of observing and remembering the flyer, r , is identified by two moments in the Opt-out treatment: the fraction of households checking the opt-out box, and the fraction opening the door. The elasticity of opening the door η with respect to incentives is identified by the fraction opening the door in the survey treatments for different payments and survey durations. In addition, η is identified by the amounts given in the different charity treatments.

The survey parameters are identified using the survey moments. The survey completion rates for varying amounts of compensation identify the heterogeneity in the willingness to complete the survey, and hence σ^{sv} . The survey completion

⁵For the results presented here, the best estimate is achieved in about 18 percent of all runs.

rate also identifies the mean willingness to complete a 10-minute survey, μ^{sv} . The value of time v^{sv} is identified from the comparison between pay increases for the survey (from \$0 to \$5 to \$10) and duration decreases (from 10 to 5 minutes). Finally, the social pressure S^{sv} is identified by the share of people answering the door in the survey treatments. To see this, consider a respondent who dislikes answering a survey and hence will say no and incur the social pressure cost S^{sv} . In the flyer treatment, she will choose to be at home with probability $h_0 - \eta r S^{sv}$ (barring corner solutions for h). Hence, knowing h_0 , η , and r , it is possible to identify S^{sv} .

Turning to the charity parameters, the information on the amounts given identify the standard deviation of altruism σ_a^{ch} , mean altruism μ_a^{ch} , and the curvature parameter Γ . This is clearest for donations of $g > g^S$, where social pressure plays no role. Without social pressure, an individual with altruism a will give exactly g dollars if the marginal utility of giving, $av'(g) = a/(\Gamma + g)$, equals the private marginal utility of consumption, 1, and hence $a = \Gamma + g$. Thus, in this example without social pressure, the mass of households with altruism higher than $\Gamma + g$, i.e., $1 - F(g + \Gamma)$, has to equal the observed share of households that give at least g . This pins down the empirical distribution of a for a given Γ . The identification of Γ depends on two sets of moments: the sorting in of givers of larger amounts, and the giving of smaller amount. The more concave the altruism function is (that is, smaller Γ), the more altruistic individuals sort in because of higher infra-marginal utility of giving, and the more frequent are small donations. Finally, the social pressure S^{ch} is identified from two main sources of variation: home presence in the flyer treatment (which, to a first approximation, equals $h_0 - \eta r S$) and the distribution of small giving (the higher the social pressure, the more likely is small giving and in particular bunching at g^S).

Estimates. Table 1 reports the benchmark estimates of the parameters along with standard errors. The probability of being at home h_0 is precisely estimated to be 39.2 percent in 2008 and 38.6 percent in 2009. The share r of households that have read (and remember) the flyer is precisely estimated at 34.6 percent. While this estimate may appear low, many households may have just disregarded the flyer, or another household member may have seen it, but not informed the person opening the door. The elasticity of home presence η is estimated to be 0.034 (s.e. 0.008), implying that the cost of increasing the probability of being at home and answering the door by 10 percentage points is $0.1^2/2\eta = \$0.15$.

We find that women and men have a very similar mean altruism for the first charity, La Rabida ($\mu = -12.26$ for women vs. $\mu = -11.35$ for men). However, women are substantially more altruistic on average towards the second charity, ECU ($\mu = -10.29$ (s.e. 1.31) for women vs. $\mu = -22.42$ (s.e. 2.01) for men), with the difference statistically significant at conventional levels.

We estimate that women have a lower variance in their altruism distribution ($\sigma^{char} = 17.42$ for women vs. $\sigma^{char} = 19.54$ for men), with the difference statistically significant due to the highly precise estimates (p=0.02).

The social pressure parameters are also quite precisely estimated. For women, turning down a door-to-door giving request is associated with a social pressure cost of \$5.01 (s.e. \$0.39) for La Rabida and \$1.28 (s.e. \$0.68) for ECU. For men, the corresponding estimates are \$3.02 (s.e. \$0.34) for La Rabida and \$2.38 (s.e. \$1.31) for ECU. Thus, we do not see a systematic relationship between gender and social pressure: women experience higher social pressure when faced with solicitors for La Rabida (the children’s hospital), while men face higher pressure when they encounter solicitors for ECU (the out-of-state research facility).

Finally, we construct a measure of “marginality” - the probability distribu-

tion function of the altruism distribution, evaluated at the threshold of altruism above which the solicitee chooses to donate a positive amount, if asked. We find that women have insignificantly higher marginality for La Rabida than men (0.014 vs. 0.013), but substantially higher marginality for ECU (0.013 vs 0.006), with the latter difference being highly statistically significant.

Turning to the survey estimates, the average utility for survey completion is estimated to equal $-\$23.57$ for women and $-\$31.96$ for men, although the difference between the two is not statistically significant. There is significant heterogeneity in survey completion utility for both genders ($\hat{\sigma}^{sv} = \$26.36$ for women and $\hat{\sigma}^{sv} = \$34.01$ for men). While the difference in standard deviations is not statistically significant by itself, the point estimates are consistent with men having greater variance in their utility for a doing a pro-social task (completing an unpaid survey for a researcher). The value of time for one hour of survey completion is imprecisely estimated to be $\$124.10$, indicative of the wealthy neighborhoods we reached.⁶ The social pressure cost of turning down a survey request, S^{sv} , is estimated to be $\$4.25$ for women and $\$10.49$ for men, sizable magnitudes. Interestingly, men are estimated to incur significantly *greater* social pressure than women when faced with a surveyor. Finally, we again show a measure of estimated “marginality” of men and women: the probability distribution function of the utility of completing the unpaid 10 minute survey, evaluated at the threshold $\bar{sv}^{m,t}$ above which the solicitee agrees to complete the survey, when asked. We estimate a higher marginality for women than men (0.012 vs 0.010), but the difference is not statistically significant.

⁶At an average income of about $\$100,000$ per year, the implied hourly wage is $\$50$.

Figure 1. Frequency of Completing the Unpaid Survey: Male versus Female




Table 1. Minimum-Distance Estimates

Common Parameters	Females		Males	
Share of Female Households	0.481 (0.006)			
Prob. of Home Presence (<i>h</i>) - Year 2008	0.392 (0.005)			
Prob. of Home Presence (<i>h</i>) - Year 2009	0.386 (0.007)			
Prob. of Observing Flyer (<i>r</i>)	0.346 (0.016)			
Elasticity of Home Presence (<i>eta</i>)	0.034 (0.008)			
Implied Cost of Altering Prob. Home by 10 pp.	0.147			
Survey Parameters				
Mean Utility (in \$) of Doing 10-Minute Survey	-23.573 (4.438)		-31.961 (6.17)	
Std. Dev. of Utility of Doing Survey	26.356 (5.971)		34.007 (8.493)	
<i>Test of equality across gender: Svy Std Dev (p value)</i>			0.440	
Social Pressure Cost if Saying No to Survey	4.255 (1.303)		10.491 (2.347)	
Value of Time of One-Hour Survey		124.100 (46.998)		
"Marginality" (pdf at altruism threshold for completing survey)	0.012 (0.003)		0.010 (0.002)	
<i>Test of equality across gender: Svy "Marginality" (p value)</i>			0.580	
Charity Parameters				
	La Rabida	ECU	La Rabida	ECU
Mean of altruism distribution, μ	-12.265 (0.996)	-10.292 (1.31)	-11.351 (0.951)	-22.419 (2.011)
Standard deviation of altruism, σ	17.422 (0.764)		19.540 (0.746)	
<i>Test of equality across gender: Std. Dev. Of Altruism (p value)</i>			0.04 **	
Social Pressure Cost of Giving 0 in Person	5.016 (0.394)	1.286 (0.686)	3.020 (0.336)	2.380 (1.31)
Curvature of Altruism Function, Γ	10.000			
"Marginality" (pdf at altruism threshold for donating)	0.014 (0.001)	0.013 (0.001)	0.013 (0.001)	0.006 (0.001)
<i>Test of equality across gender: "Marginality" (p value)</i>	0.48	0.00 ***		
SSE	178.330			

Notes: Estimates from minimum-distance estimator with weights given by inverse of diagonal of variance-covariance matrix. Standard errors are in parentheses. SSE reports the Weighted Sum of Squared Errors. "Marginality" for survey is reported for an unpaid survey of length 10 min. A solicitee will be willing to do such a survey if the utility of doing the survey exceeds the social pressure cost of saying no, i.e. $s > -S^s$. The marginality is thus defined as $f(-S^s, \mu^s, \sigma^s)$, where $f(x; \mu, \sigma)$ is the pdf of the normal distribution with mean μ and variance σ^2 . Similarly, DellaVigna et al (2012) show that a solicitee will donate a positive amount of money if she has altruism $a > (1-S^{char}) * \Gamma$. Marginality is thus measured as $f((1-S^{char}) * \Gamma, \mu^{char}, \sigma^{char})$, where $char = \{la\ rabida, ecu\}$ and $f(\cdot)$ is again the pdf of the normal distribution.