

“LAST-PLACE AVERSION”: EVIDENCE AND REDISTRIBUTIVE IMPLICATIONS*

ILYANA KUZIEMKO
RYAN W. BUELL
TALY REICH
MICHAEL I. NORTON

We present evidence from laboratory experiments showing that individuals are “last-place averse.” Participants choose gambles with the potential to move them out of last place that they reject when randomly placed in other parts of the distribution. In modified dictator games, participants randomly placed in second-to-last place are the most likely to give money to the person one rank above them instead of the person one rank below. Last-place aversion suggests that low-income individuals might oppose redistribution because it could differentially help the group just beneath them. Using survey data, we show that individuals making just above the minimum wage are the most likely to oppose its increase. Similarly, in the General Social Survey, those above poverty but below median income support redistribution significantly less than their background characteristics would predict. *JEL* Codes: H23, D31, C91.

I. INTRODUCTION

A large literature in economics argues that utility is related not only to absolute consumption or wealth but also to an individual’s relative position along these dimensions within a given reference group.¹ This literature has shown that both ordinal and cardinal comparisons affect utility, but it has devoted less attention to the shape of these effects over the distribution. This article reports results from laboratory experiments designed to test whether ordinal rank matters differently to individuals depending on their position in the distribution. We hypothesize that individuals exhibit a particular aversion to being in last place, such

*We thank Roland Bénabou, Doug Bernheim, Angus Deaton, Ray Fisman, Claudia Goldin, Caroline Hoxby, Seema Jayachandran, Emir Kamenica, Judd Kessler, David Moss, Sendhil Mullainathan, Shanna Rose, Al Roth, Stefanie Stantcheva, Gui Woolston, and participants at seminars at NBER, Brown, Columbia, Michigan, Ohio State, Princeton, and Stanford for helpful comments and conversations. We are indebted to Larry Katz, Andrei Shleifer, and four anonymous referees for suggestions that greatly improved the quality of the article. Jimmy Charite, Elena Nikolova, and Sam Tang provided excellent research assistance.

1. This work largely began with Duesenberry (1949). We review the literature more thoroughly later in the Introduction.

© The Author(s) 2013. Published by Oxford University Press, on behalf of President and Fellows of Harvard College. All rights reserved. For Permissions, please email: journals.permissions@oup.com

The Quarterly Journal of Economics (2014), 105–149. doi:10.1093/qje/qjt035.

Advance Access publication on November 13, 2013.

that a potential drop in rank creates the greatest disutility for those already near the bottom of the distribution.

Our second objective is to explore how this “last-place aversion” predicts individuals’ redistributive preferences outside the laboratory. Many scholars have asked why low-income individuals often oppose redistributive policies that would seem to be in their economic interest. Last-place aversion suggests that low-income individuals might oppose redistribution because they fear it might differentially help a last-place group to whom they can currently feel superior. We present supporting evidence for this idea from survey data, though identifying last-place aversion outside the laboratory is admittedly more challenging because it is relatively harder to determine where individuals see themselves in the income distribution.

We begin with the more straightforward task of identifying last-place aversion (LPA) in laboratory experiments. The two sets of experiments explore LPA in very different contexts. In the first set of experiments, participants are randomly given unique dollar amounts and then shown the resulting “wealth” distribution. Each player is then given the choice between receiving a payment with probability 1 and playing a two-outcome lottery of equivalent expected value, where the “winning” outcome allows the possibility of moving up in rank. We find that the probability of choosing the lottery is uniform across the distribution except for the last-place player, who chooses the lottery significantly more often.

In the second set of experiments, participants play modified dictator games. Individuals are randomly assigned a unique dollar amount, with each player separated by a single dollar, and then shown the resulting distribution. They are then given an additional \$2, which they must give to the person either directly above or below them in the distribution. Giving the \$2 to the person below means that the individual herself will fall in rank, as ranks are separated by \$1. Nonetheless, players almost always choose to give the money to the person below them, consistent with Fehr and Schmidt (1999) and other work on inequality aversion. However, the subject in second-to-last place gives the money to the person *above* her between one-half and one-fourth of the time, consistent with LPA’s prediction that concern about relative status will be greatest for individuals who are at risk of falling into last place.

In our data it is sometimes difficult to distinguish between strict last-place aversion and a more general low-rank aversion—in some experiments, individuals seem modestly averse to second-to-last place as well as last place. However, we can always separate last-place or low-rank aversion from a desire to be above the median, the inequality-aversion model of Fehr and Schmidt (1999), the equity-reciprocity-competition model of Bolton and Ockenfels (2000), the distributional preference model of Charness and Rabin (2002), total-surplus maximization, and, generally, a linear effect of initial rank.

Although supportive of LPA, the laboratory evidence alone is limited because it can only speak to whether the phenomenon exists in game-like settings. We thus turn to survey data to examine whether patterns of support for actual redistributive policies are consistent with LPA. Of course, in the real world, the concept of last place is far less well defined than in the two experimental environments described here. To a first approximation, no one is literally in last place in the U.S. income or wealth distribution. Strictly speaking, LPA cannot explain why, for example, politicians might be able to divide low-income voters and prevent them from uniting in support of redistributive taxes and transfers, because such policies will not land anyone literally in last place. If, instead, individuals create reference groups specific to whatever policy question they are considering, then LPA has more hope of explaining policy preferences.

We begin with the minimum wage. LPA predicts that those making just above the current minimum wage might actually oppose an increase—though they might see a small increase in their own wage, they would now have the last-place wage themselves and would no longer have a group of worse-off workers from whom they could readily distinguish themselves. We could not find existing survey data that includes both respondents' actual wages (as opposed to family income) and their opinion regarding minimum wage increases, so we conducted our own survey of low-wage workers. Consistent with almost all past surveys on the minimum wage, support for an increase is generally over 80%. However, consistent with LPA, support for an increase among those making between \$7.26 and \$8.25 (that is, within \$1 more than the current minimum wage of \$7.25 and thus those most likely to “drop” into last place) is significantly lower.

Finally, we use nationally representative survey data to examine whether more general redistributive preferences

appear consistent with LPA. In particular, do those who are above poverty but below median income—roughly speaking, the analogue to the second-to-last-place subjects in our redistribution experiment—exhibit softer support for redistribution than their background characteristics would otherwise suggest? Although hardly a definitive test of LPA, we would be concerned if individuals relatively close to the bottom of the distribution were highly supportive of redistribution. In fact, in General Social Survey data, the pattern predicted by LPA holds across a variety of survey questions and subgroups of the population.

Our article contributes to the literature on distributional preferences, which many past authors have explored using, as we do, modified dictator games. As most of these experiments involve just two (or at most three) players, they can offer only a very limited view of the shape of distributional preferences as a function of relative position.² Moreover, as we discuss, many of these models (e.g., those positing that individuals wish to improve the position of the worst-off person) have predictions in the *opposite* direction of LPA. In general, we show that many of the predictions of these models tend to break down for individuals near the bottom of the distribution.

In contrast to the experimental approach, other papers have used survey data to examine how subjective well-being varies with one's position in the income distribution, though they have not tested our specific nonlinear formulation. Boyce, Brown, and Moore (2010) use British data to show that percentile in the income distribution predicts life satisfaction better than either absolute income or relative income (absolute income divided by some reference income level, usually the mean or median). Clark, Westergård-Nielsen, and Kristensen (2009) are able to focus on small Danish neighborhoods and find that income rank within locality is a better predictor of economic satisfaction than is absolute income.

Instead of examining strictly ordinal measures like rank or percentile, most papers in this literature have instead focused on

2. As Engelmann and Strobel (2007) note in their review of distribution games: "Taking note of the limited ability of two-player dictator games to discriminate between different distributional motives . . . it is surprising that there is a relative sparsity of dictator experiments with more than two players." A recent addition to the literature is Durante, Putterman, and van der Weele (forthcoming), who conduct 20-player distribution experiments, though in most of their sessions players do not know their place in the distribution.

relative income, probably because it requires knowing only the mean or median (as opposed to the entire distribution) of the comparison group's income. Luttmer (2005) and Blanchflower and Oswald (2004) find that holding own income constant, increasing the income of those living near you has a negative effect on reported well-being; Hamermesh (1975) provides an early example of a similar effect regarding relative wages and job satisfaction and Card et al. (2012) use an experiment in which only some employees are encouraged to learn their relative wage to demonstrate the same result. There is no consensus on whether there is a nonlinear effect of relative income—Card et al. find that those below the median care more about relative income, Blanchflower and Oswald find some evidence in the opposite direction, and Luttmer finds those below and above the median are affected equally by relative income.

Other papers have focused on why individuals care about relative position. Cole, Mailath, and Postlewaite (1992) argue that even if individuals do not care about relative position per se (i.e., it does not enter directly into their utility function) because many real outcomes (such as marriage quality) depend on relative as opposed to absolute position, relative position will appear in reduced-form utility expressions. Also focusing on relative competition, Eaton and Eswaran (2003) develop a model in which natural selection favors those who care about relative position, as relative position determines access to food sources and high-quality partners. Indeed, Raleigh et al. (1984) offer empirical evidence that concern about rank is “hard-wired”—they find that when a dominant (subordinate) vervet monkey is placed in a group where he is now subordinate (dominant), his serotonin level drops (rises) by 40–50%.³

By the logic developed in these evolutionary models, not only would humans care about relative position in general but a strong aversion to being near last place would arise because in a monogamous society with roughly balanced sex ratios, only those at the very bottom would not marry or reproduce. Indeed, being “picked last in gym class” is so often described as a child's worst fear that the expression has become a cliché.

Although few papers have linked social comparison to support for redistributive policies, there is a large literature on how individuals form redistributive preferences. Many studies have

3. See Zizzo (2002) for a review of the neurobiology of relative position.

examined how demographic and background characteristics determine support for redistribution.⁴ Others have focused, as we do, on explanations for why low-income voters do not support higher levels of redistribution, examining on mobility (Benabou and Ok 2001), imperfect information (Bredemeier 2010), and the role of competing, noneconomic issues that divide low-income voters (Roemer 1998). In the U.S. context, race has often been examined as one such issue (Lee and Roemer 2006). In our analysis of the General Social Survey, we thus take care to show that redistributive preference patterns consistent with LPA hold for both whites and minorities and thus cannot be explained merely by whites' views of low-income minorities.

Although we focus on redistributive preferences and risk taking, researchers have examined other potential consequences of social comparison. For example, Veblen (1899) argued that concern for relative position inspires conspicuous consumption, an idea formalized by Frank (1985) and others, and explored empirically by Charles, Hurst, and Roussanov (2009).

The remainder of the article is organized as follows. Section II discusses how LPA can be separately identified from other models of preferences and social comparison. Sections III and IV presents results from, respectively, the lottery experiment and the modified dictator experiment. Sections V and VI include the results from our minimum wage survey and the General Social Survey analysis, respectively. Section VII discusses the potential implications of LPA for behaviors beyond those we investigate in this article and offers recommendations for future work.

II. SEPARATING LAST-PLACE AVERSION FROM OTHER MODELS OF PREFERENCES

Consider a finite number of individuals with distinct wealth levels $y_1 < y_2 < \dots < y_N$, so y_1 is the wealth of the poorest (last-place) person. We follow previous research and assume that utility is additively separable in absolute wealth and relative position.⁵ We write the utility of person i as:

$$u(y_i, r_i) = \gamma g(r_i) + (1 - \gamma)f(\cdot),$$

4. For example, see Alesina and Giuliano (2011) and citations therein.

5. For example, see Charness and Rabin (2002), but also many others.

where r_i is *relative position*, so $r_i = 1$ for the last-place person, up to $r_i = N$ for the first-place person. Let $\gamma \in (0, 1)$. For the moment, we set aside $f(\cdot)$ and focus on $g(r)$.

Strictly speaking, LPA assumes that $g(r_i) = \mathbb{1}(r_i > 1) \equiv \mathbb{1}(y_i > y_1)$, where $\mathbb{1}$ is an indicator function. Essentially, $g(r_i)$ is a bonus payment to all but the last-place individual. We plot this function for the six-person distribution $y_1 = \$1, y_2 = \$2, \dots, y_6 = \$6$ (which will be used in some of our experiments) in Figure I.

If, instead, one assumes that individuals have a special dislike for being low-rank, not just last-place, then $g(r_i)$ would not be a step function but a positive, concave function where utility increases steeply at the bottom of the distribution but then quickly flattens out. This shape also arises if $g(r_i) = \mathbb{1}(r_i > 1) \equiv \mathbb{1}(y_i > y_1)$, but y_i is subject to random perturbations ϵ_i . The second series of Figure I plots the probability that $y_i + \epsilon_i$ ($\epsilon \sim \mathcal{N}(0, 1)$) is above last place in the ex post distribution, by the original rank in the ex ante $\{y_1, y_2, \dots, y_6\}$ distribution. There is substantial disutility faced not only by the last-place player but also the second-to-last: he faces the nontrivial probability of falling into last place given the income uncertainty, whereas this risk is essentially zero for the individuals above him.

We take an agnostic view of how to specify $f(\cdot)$ and instead focus on empirically separating LPA from a large class of $f(\cdot)$ functions posed in the existing literature, so f can take a variety of arguments, such as absolute and relative levels of income as well as functions of ordinal rank besides LPA. LPA suggests that the predictions from many of these models will break down for individuals near the bottom of the distribution.

For example, as we discuss in the next section, standard formulations of expected utility theory, that is, $f(\cdot) = f(y_i)$, predict that individuals will only choose a lottery over a risk-free payment of equal expected value if they are risk-seeking. Given that risk aversion is believed to decrease in wealth (and that this prediction holds over small stakes in laboratory settings), it further predicts that the last-place player would be the *least* likely to choose the lottery. As we show in the next section, LPA predicts that last-place individuals will be the *most* likely to choose the lottery, as long as it offers them the chance to move up in rank.

Similarly, we show in Section IV that LPA can be separately identified from many distributional preference models. For

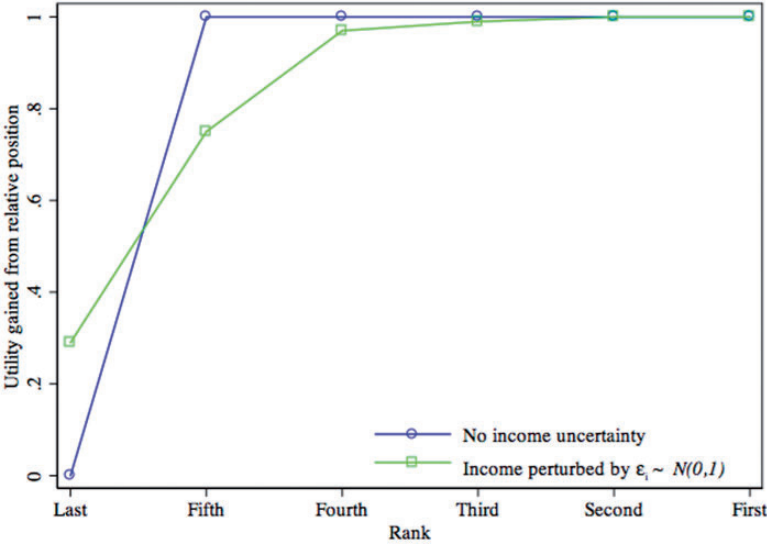


FIGURE I

Probability of Being Above Last Place in Ex Post Distribution, by Ex Ante Rank
(Ex Ante Distribution is $y_1 = \$1, \dots, y_6 = \6)

In the first series there is no income uncertainty, and thus the probability of being in last place is 1 for the current last-place player and 0 for others. The probabilities plotted in the second series are generated as follows. We begin with the ex ante income distribution $y_1 = \$1, y_2 = \$2, \dots, y_6 = \$6$. We transform it into the ex post distribution by adding an independently drawn ϵ_i to each y_i , where $\epsilon_i \sim \mathcal{N}(0, 1)$. After these six draws, the individuals are reranked based on the ex post distribution. We repeat the process 10,000 times. The probability of being above last place in the ex post distribution for each ex ante rank was averaged over the 10,000 repetitions.

example, in their model of inequality aversion, Fehr and Schmidt (1999) posit that

$$\begin{aligned}
 f(\cdot) &\equiv f(y_1, y_2, \dots, y_i, \dots, y_N) \\
 &= y_i - \alpha \frac{\sum_{j \neq i} \max\{x_j - x_i, 0\}}{n - 1} - \beta \frac{\sum_{j \neq i} \max\{x_i - x_j, 0\}}{n - 1},
 \end{aligned}$$

where $\alpha > \beta > 0$. This model predicts that if a person is given a choice between giving money between a person above him in the distribution or a person below him, he will always choose the person below him. LPA suggests that this prediction will break

down for individuals just above last place, because giving to the person below moves them closer to last place themselves.

In the sections that follow, we empirically separate the predictions of LPA from these and other models posited by the existing literature.

III. EXPERIMENTAL EVIDENCE OF LAST-PLACE AVERSION: MAKING RISKY CHOICES

We begin our test of LPA by examining whether individuals choose to bear risk in return for the possibility of moving out of last place. Note that we made an effort to design these experiments as well as those in Section IV so as not to unduly trigger LPA. As we speculate that shame or embarrassment may motivate individuals' desire to avoid last place, participants never interact face to face, but through computers, and they generate their own screen names and are thus free to protect their identity if they wish.⁶ We seated everyone walking into the lab sequentially into different rows, so those who entered together and presumably might know each other were not in the same group and thus individuals did not play against their friends. Each individual sits in a separate carrel surrounded by large blinders, which further enhance privacy and anonymity. Players are not publicly paid; instead money is discreetly given to them while they are still sitting in their carrels.⁷ Finally, all of the experiments involve an initial assignment to a rank, and we make clear to participants that this assignment is performed randomly by a computer. We believe the emphasis on random assignment should diminish LPA by discouraging players from associating rank and merit.

Despite these steps, it is also a fair critique that in attempting to make the experiment engaging, we may have put participants in the mindset of playing a game (readers can judge for

6. For confidentiality reasons, the data extract that we receive from the lab does not include respondents' actual names, so we cannot examine how many chose "fake" screen names. But just under 20% use screen names that are obviously not real names (e.g., "turtle," "panda," "Big Papi"). Moreover, in almost no cases did people use screen names that looked like a first and a last name, so subjects would be unable to look up their opponents on, say, Facebook or Google after the experiment.

7. Although privacy likely diminishes LPA, it is unlikely to eliminate it—the literature discussed in the Introduction suggests that concern for rank may be hard-wired and in fact field work suggests that effort changes when individuals learn their rank privately (see, e.g., Tran and Zeckhauser 2012).

themselves by examining the screenshots in the Online Appendix). Finishing in first or last place may well be especially salient in games, and as such this design may unfairly trigger LPA. On the other hand, as the screenshots show, when we display ranks, we always describe the last-place player as being in N th place in a N -player game, never *in last place*. Moreover, in the real world one's economic status is neither private nor explicitly random, suggesting that our experimental design may if anything diminish status concerns relative to those experienced outside the lab.

III.A. Data and Experimental Design, Main Experiment

Participants ($N=84$) sign up by registering online at the Harvard Business School Computer Lab for Experimental Research (CLER). See Online Appendix Table 1 for demographic summary statistics as well as more detailed information on eligibility requirements for registration and payment of participants. We randomly divide participants into 14 groups of 6, with groups being fixed across all rounds. Each round begins with the computer randomly assigning each player a place in the distribution $\{\$1.75, \$2.00, \$2.25, \$2.50, \$2.75, \$3.00\}$. Ranks and actual dollar amounts of all players are common knowledge and clearly displayed throughout the game. The computer then presents an identical two-option choice set to all players in the game:

In this round, which would you prefer?

- (i) Win \$0.13 with 100% probability.
- (ii) Win \$0.50 with 75% probability and lose \$1.00 with 25% probability.

After players have submitted their choices, the computer makes independent draws from the common $P(\text{win}) = 3/4$ probability distribution for each player who chooses the lottery and adds the risk-free amount to the balance of each player who did not choose the lottery. The new balances and ranks are then displayed. The players are then re-randomized to the same $\{\$1.75, \dots, \$3.00\}$ distribution and the game repeats. Each session consists of nine rounds, but participants are not told how many rounds the game entails to avoid end effects.⁸ Participants are told that one

8. See Rapoport and Dale (1966) for an early treatment of so-called end effects.

randomly selected player from the session will be paid his balance from one randomly selected round.

Note that the payment players can receive with probability 1 is always equal to half the difference between ranks, rounded up to the nearest penny. That is, $\$0.125$ ($\$0.25 \div 2 = \0.125), rounded up to $\$0.13$. The “winning” payment of the lottery is always equal to the difference between a given individual and the person two ranks above him, that is, $\$0.50$. The losing outcome of the lottery is set to $\$1$, so that the lottery and the risk-free options are equal in expected value after rounding ($0.75 \cdot 0.50 - 0.25 \cdot 1 = 0.125 \approx 0.13$), and for ease of exposition we describe the two options as having equal expected value. Note that even if the last-place player chooses the lottery and loses, he will still have $\$0.75$, so can never “owe” money.

III.B. Predictions

What does existing literature suggest about laboratory participants’ tendencies to choose the lottery over the risk-free option? All research that we found assessing risk aversion in the lab explores settings without social comparison (e.g., individuals do not interact with others and only know their own experimental income levels). First, in contrast to strict risk aversion, existing work suggests that between one-fourth and one-half of subjects appear risk-seeking or risk-neutral in laboratory experiments.⁹ Second, past work suggests that any such risk taking should increase with initial wealth levels. That is, in the laboratory subjects display diminishing absolute risk aversion.¹⁰

LPA offers predictions that are in sharp contrast to diminishing absolute risk aversion. As we show later, the exact predictions depend slightly on players’ levels of strategic

9. Holt and Laury (2002) find that subjects choose the riskier of two options about one-third of the time. Harrison, List and Towe (2007) use a similar procedure and find that 56% in fact choose the riskier lottery. Dohmen et al. (2005) find that roughly 22% are risk neutral or risk loving, even in situations with relatively large stakes. In perhaps the application closest to ours, in that subjects choose between lotteries and risk-free payments of equal expected value, Harbaugh, Krause, and Vesterlund (2002) find that 46% of adult laboratory subjects choose the lottery. Note that we do not cite evidence on risk aversion outside the laboratory, given the critique that the risk aversion displayed over small stakes in laboratory settings is perhaps a separate phenomenon from that displayed in the real world (Rabin 2000).

10. See Levy (1994), Holt and Laury (2002), and Heinemann (2008), among many others.

sophistication—because players make their choices simultaneously, more strategically minded players would condition their choice on what they think others will do. However, under all sophistication assumptions, we predict that those in the *bottom* of the distribution will choose the lottery more often than those at the top.

First, assume that, as a heuristic, players hold others' balances constant when they make their decisions.¹¹ LPA then predicts that the last-place player will choose the lottery more often than other players will. We relegate the algebra to the Online Appendix, but the intuition is simple: only for the last-place player does the lottery offer a chance to move out of last place, and thus even some risk-averse subjects will find that this possibility outweighs the utility cost of bearing additional risk.

Second, assume that instead of holding others' balances constant, individuals assume that their fellow subjects choose randomly between the lottery and risk-free option (that is, they assume their fellow players are level-0 reasoners, meaning they themselves are level-1 reasoners, to borrow the terminology in Stahl and Wilson 1995).¹² To predict decisions under this set of assumptions, for each player, we simulate the resulting distribution when (i) he chooses the risk-free option, versus (ii) he chooses the lottery, where in both cases his fellow subjects play randomly. As Online Appendix Figure 1 shows, the probability of escaping last place is maximized for the last-place player when he chooses the lottery, but for all other players the probability of avoiding last place is maximized by taking the risk-free option. As such, LPA again predicts that the last-place player will be the most likely to choose the lottery.

Finally, players may assume their opponents play strategically and thus solve for the Nash equilibrium. In Online Appendix D, we show that under LPA, the incentive to gain the $\gamma^l(y_i > y_1)$ term of the utility function has the last- and second-to-last-place players playing a mixed strategy between the lottery and risk-free option, with no one else choosing the lottery. Assuming again

11. For example, see Moore, Oesch, and Zietsma (2007), and Radzevick and Moore (2008), who find that subjects ignore their opponents' decisions even in situations where those decisions should be highly salient.

12. A level-0 reasoner makes decisions randomly. In the Stahl and Wilson (1995) terminology, a level- k reasoner assumes that his opponents are drawn from a distribution of level-0 through level- $k-1$ reasoners. As such, a level-1 reasoner assumes that his opponents play randomly.

that there is some baseline level of risk-seeking subjects in our laboratory settings, LPA under this scenario predicts that the last- and second-to-last-place subjects will choose the lottery at a greater tendency than will other subjects.

Subjects' strategic sophistication is difficult to predict a priori. Much work has found that subjects display level-1 sophistication, suggesting we would see elevated risk taking for the last-place player but not the second-to-last.¹³ Moreover, it has been shown that subjects are less likely to converge to Nash play when the Nash equilibrium is in mixed strategies.¹⁴ Either way, the prediction from LPA contrasts to that of the standard model and thus the experiment provides a demanding test of our theory.

III.C. Results

Figure II shows the probability individuals choose to play the lottery, as a function of their rank at the time they make the decision. Its most striking feature is the relatively flat relationship between rank and the propensity to choose the lottery for ranks one through five, contrasted with the elevated propensity for players in last place. As the regression analysis will show, not only is the last-place player significantly more likely than other players to choose the lottery, the p -values noted on the figure show that the pair-wise difference between the last-place player and each of the other individual ranks are generally statistically significant (with the one exception having a p -value of 0.128). Online Appendix Figure 2 shows that the elevated tendency of the last-place subject to choose the lottery holds after excluding the first two rounds, which previous research has shown are noisier as players are still learning.¹⁵

Although the last-place player plays the lottery most often, other players do not completely eschew it. This finding is consistent with the literature cited earlier on risk-seeking behavior in laboratory settings, though the rates of risk seeking for ranks one through five appear somewhat higher in our experiment. There is no evidence that the fifth-place player is "defending" his position

13. For example, see Nagel (1995) and Costa-Gomes and Weizsäcker (2008), in which most players appear to be level 1 and Camerer, Ho, and Chong (2004), where players appear to be level 1.5.

14. See Ochs (1995).

15. See Carlsson (2010) for a discussion and review of literature on why preferences may be more stable as subjects gain experience.

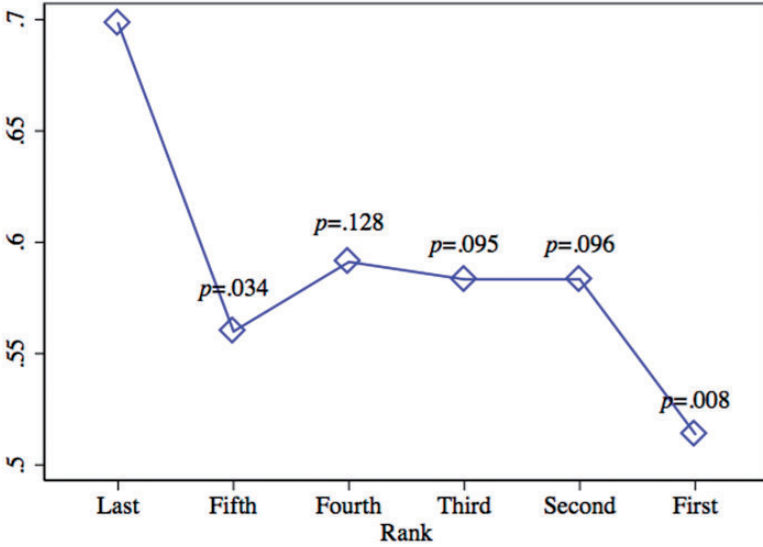


FIGURE II

Probability of Choosing the Lottery Over the Risk-Free Payment (One-Shot Games)

Based on 14 six-player games of nine rounds each, for a total of 756 observations. Each round every player was given the same choice between a two-outcome lottery and a risk-free payment of equivalent expected value. See Section III.A for details. All coefficients and *p*-values are based on the probit regression: $chouselottery_i = \sum_{k=1}^5 \beta^k rank_i^k + \epsilon_i$, where $rank_i^k$ is an indicator variable for player *i* having rank *k*, standard errors are clustered by player, and no other controls are included. As there are six ranks in the lottery experiment, the excluded category is sixth (last) place. The *p*-values in the figure refer to the estimated fixed effect of being in rank *k* relative to the excluded category of last place. The *y*-axis values are the probit coefficients (as changes in probability) plus the mean of the excluded category (to normalize).

against the heightened tendency of the last-place player to gamble, as in the Nash outcome. Note also that there is no evidence of a first-place effect—those in first and second place do not appear to compete against each other by going for the higher potential payoff. If LPA were purely being driven by the game-like setting of the experiment, one might also expect a competitive effect at the top of the distribution, given the salience of “first place” in games.

Table I displays results from probit regressions, reported as changes in probability. Column (1) includes dummies for fifth and sixth place as well as round and group fixed effects. The results

TABLE I
 PROBIT REGRESSIONS OF THE PROPENSITY TO CHOOSE THE LOTTERY OVER THE RISK-FREE PAYMENT

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
In last place	0.134** [0.0533]	0.135** [0.0518]	0.128** [0.0619]	0.125** [0.0564]	0.110** [0.0512]	0.109* [0.0571]	0.179* [0.0914]	0.104 [0.0639]
In fifth place	-0.00600 [0.0476]							
Below median				0.0171 [0.0378]				
Below previous round					0.0583 [0.0412]			
Δ Exp. rank						0.0225 [0.0216]		
Δ Disadv. inequality							-0.0415 [0.0491]	
Δ Adv. inequality							0.322 [0.237]	
Rank, scaled 0 to 1								0.0528 [0.0648]
Mean, dept. var.	0.591 All	0.591 All	0.592 Ex. early	0.591 All	0.591 All	0.591 All	0.591 All	0.591 All
Log likelihood	-475.7	-475.7	-371.6	-475.6	-474.6	-475.1	-474.5	-475.3
Observations	756	756	588	756	756	756	756	756

Notes. All regressions are estimated via probit, coefficients are reported as marginal changes in probability, and standard errors are clustered by individual. The sample is based on 14 six-player games of nine rounds each. The dependent variable for all regressions is an indicator variable coded as 1 if the subject chose the lottery over the risk-free payment. See Section III for further details on the experiment. In specifications that “exclude early” rounds, the first two rounds are not included. “Below median” is an indicator for whether the individual was in the bottom half of the distribution at the time of his decision. “Below previous round” is an indicator for beginning this round with a lower rank than the previous round (it is coded as 0 for the first round). $\Delta Exp. Rank$ is defined as the expected change in rank from playing the lottery (holding other balances constant). Following Fehr and Schmidt (1999), *Disadvantageous inequality* is defined as $\sum_{j \neq i} \max\{y_j - x_i, 0\}$ and *Advantageous inequality* as $\sum_{j \neq i} \max\{x_i - y_j, 0\}$. The Δ for each of these variables is defined as the expected value when player i plays the lottery minus the value when he takes the risk-free payment (holding other balances constant). “Rank, scaled 0 to 1” is an individual’s rank at the time of his decision, scaled so that first place is 0 and last place is 1. * $p < .10$, ** $p < .05$, *** $p < .01$.

suggest that last-place players play the lottery 13.4 percentage points (or 23%, given a mean rate of playing the lottery of 0.591) more than players in ranks one through four, and there is no differential effect for the fifth-place player. For the rest of the table we pool the fifth-place player with ranks one through four to gain power. Columns (2) and (3) show that the last-place effect is robust to excluding the first two rounds.

Prospect theory suggests that individuals are risk-loving over losses when they are below their reference point. Two of the most commonly posited reference points are the group mean or median (which in our case are equivalent) and one's previous outcome.¹⁶ Columns (4) and (5) show results when, respectively, we add controls for being below the median (in which case LPA is identified by comparing last place to fourth and fifth place) or below one's previous outcome. In both cases these controls have the expected, positive sign, but their inclusion does not affect the coefficient on the last-place term.

A potential confound in the experiment is that those at the bottom (top) of the distribution have only limited ability to fall (rise) in rank, which might increase (decrease) their incentive to gamble. In column (6) we thus control for each player's expected change in rank from choosing the lottery over the risk-free payment.¹⁷ The coefficient on this term has the expected, positive sign, but the coefficient on the last-place term remains positive and significant.

While inequality-aversion has not often been applied to decisions over risk, we explore this possibility in column (7). We calculate the expected value of the two Fehr-Schmidt terms under two scenarios: (i) player i plays the lottery and all other players' balances are held constant; (ii) player i takes the safe option, and all other players' balances are held constant.¹⁸ For each player,

16. See Clark, Frijters and Shields (2008) for a review of the literature on how individuals form reference points.

17. For example, for $rank = 2$, winning \$0.50 would increase his rank by 1, whereas losing \$1 would decrease his rank by 4, so the expected change in rank from playing the lottery is $0.75 \cdot 1 - 0.25 \cdot 4 = -0.25$. By construction, choosing the risk-free payment does not change his rank. As such, the expected change in rank from playing the lottery relative to taking the risk-free payment, $\Delta Exp. rank$, is $-0.25 - 0 = -0.25$.

18. In Online Appendix Table 2 we show that results are robust if instead we do these calculations assuming that player i chooses the lottery and wins or that player i chooses the lottery and loses.

we take the difference in disadvantageous (advantageous) inequality under these two scenarios as a proxy for the net effect of his decision on disadvantageous (advantageous) inequality. The results in column (7) suggest that adding these controls increases the propensity of the last-place player to play the lottery.

In column (8) we evaluate LPA versus a model where the effect of rank is linear. Adding a linear rank control (for ease of interpretation, we scale it so that first place is coded as 0 and last place is coded as 1) has only a marginal effect on the last-place dummy (it falls by less than one-fourth from its value in column (2)) though the p -value is now .108 as standard errors have increased. The coefficient on linear rank is small—moving from first to last place increases the propensity to choose the lottery by only 0.05 percentage point, whereas the last-place effect by itself is equal to 0.104 percentage point.

Online Appendix Table 3 shows that the main results in column (2) of Table I are robust to adding background and demographic controls. The only significant differential LPA effect we find is that men are more likely than women to choose the lottery when in last place. The table also shows that the results barely change when individual fixed effects are included. We tend not to emphasize these results, as with only nine rounds there is still considerable between-player variation in the randomly assigned ranks that is useful to exploit.

III.D. Last-Place Aversion When Balances Accumulate

It is possible that rerandomizing ranks each round increases the propensity of low-rank players to choose the lottery, as the consequences are limited to the given round and there is no risk of accumulating large losses. In an additional experiment, we let balances accumulate between rounds to test the robustness of the LPA effect.¹⁹

The “risk-free” and “lottery” options of the first round of this experiment are equivalent to that of the original experiment. However, unlike the first experiment, players’ balances evolve after the first round and thus we modify the values of the risk-free and lottery options accordingly. The risk-free payment is always equal to half the difference between the current balance

19. This experiment as well as the two described in the next section took place in separate sessions (for a total of four sessions), so subjects are not contaminated across experiments.

of the last- and fifth-place players. The “winning” payment of the lottery is always equal to the difference between the last- and fourth-place players. As before, the payoffs are designed so that last-place individuals always have the opportunity to accept a gamble that offers the possibility of moving up in rank, holding all other players’ balances constant; this condition holds 92% of the time for ranks two through five as well. Because balances accumulate, players who lose successive lotteries can have negative balances.²⁰ The notes to Online Appendix Table 4 offer further detail.

This version of the game has the drawback that incentives are more difficult to model in a dynamic setting than in a one-shot game: as players are paid based on a randomly chosen round, in principle they should weigh both the immediate effect of their decision (equivalent to the one-shot game) as well as the effects on later rounds. However, given the evidence suggesting that subjects tend to maximize current-round payoffs even in multi-round games where the actual payoff is explicitly based on the final balance, it is likely that subjects will generally think of their decisions as in one-shot games.²¹ Despite this ambiguity, this experimental design has the important benefit that accumulating balances better reflect the real world, where income and wealth are not rerandomized at the start of each period.

Figure III shows that the heightened tendency of the last-place player to choose the lottery is not merely an artifact of rerandomization and remains when balances accumulate. In this case, we find evidence that the second-to-last-place player chooses the lottery at a significantly higher rate as well. In fact, he is marginally more likely to choose the lottery than the last-place player, though this difference is reduced when we drop the first two rounds and disappears when we control for whether losing the lottery would lead to a negative balance.²² Online Appendix Table 4 shows that the heightened tendency of the fifth- and sixth-place players to choose the lottery is robust to the alternative hypotheses we explored in Table I (in particular, the effect remains highly significant after a linear rank term is

20. In this version of the game, we give a \$20 bonus payment to the player randomly chosen to receive his experimental earnings, so that players never owe money.

21. For example, see Benartzi and Thaler (1995), Gneezy and Potters (1997), and Camerer (2003) among others.

22. See Online Appendix Figure 3 and Online Appendix Table 5.




FIGURE III

Probability of Choosing the Lottery Over the Risk-Free Payment When Balances Accumulate

Based on 12 six-player games of nine rounds each, for a total of 648 observations. Each round every player was given the same choice between a two-outcome lottery and a risk-free payment of equivalent expected value. See Section III.D for details. All p -values are based on the probit regression: $chotelottery_i = \sum_{k=1}^4 \beta^k rank_i^k + \epsilon_i$, where $rank_i^k$ is an indicator variable for player i having rank k , standard errors are clustered by player, and no other controls are included. As there are six ranks in the lottery experiment, the excluded category is being in sixth (last) or fifth place. The p -values in the figure refer to the estimated fixed effect of being in rank k relative to the excluded category of last or fifth place. The y -axis values are the probit coefficients (as changes in probability) when $chotelottery$ is regressed on ranks one through five plus the mean for the last-place player (to normalize).

included). A nice feature of this version of the experiment is that in some rounds, even if he wins the lottery the last-place player cannot “catch” the fifth-place player as long as the fifth-place player takes the risk-free payment. Consistent with LPA, there is no heightened tendency for lower-rank players to choose the lottery in these rounds.

Given recent work showing that highly salient feedback and fixed-partner matching enhances strategic play, we suspect that this dynamic form of the game makes fifth-place players move

toward the Nash outcome of the stage game.²³ Feedback is obviously more salient in this version of the game, as your current decision affects your future balance. Because ranks are “stickier” in this version of the game, players are much more likely to have the same person one rank above them two rounds in a row, thus approximating “fixed partner” matching in two-person games.²⁴

The results in this section contrast sharply with previous experimental findings—which come from settings without social comparison—showing that subjects exhibit diminishing absolute risk aversion in the lab. As already noted, our subjects also exhibit somewhat higher levels of risk-seeking. Our results thus suggest that both the level of risk aversion and its relationship to experimental wealth may depend on whether individuals view wealth in an absolute or relative sense, an interesting question for future research.

IV. EXPERIMENTAL EVIDENCE OF LAST-PLACE AVERSION: PREFERENCES OVER REDISTRIBUTION

In this section, we test the predictions of LPA in a very different context—individuals’ decisions to redistribute experimental earnings among their fellow players. Following previous research, we explore distributional preferences using modified dictator games. However, unlike previous research, which is generally restricted to experiments with two or at most three players, we examine a large enough distribution to meaningfully explore nonlinearities in preferences with respect to relative position.

IV.A. *Experimental Design*

As in the lottery experiment, the game begins with players ($N = 42$, divided into seven six-player games) being randomly assigned dollar amounts, in this case \$1, \$2, . . . , \$6. As before, the ranks and current balances of all players are common knowledge throughout the game. Each player ranked two through five must choose between giving the player directly above or directly below them an additional \$2. As players are separated by \$1, giving to

23. See Hyndman et al. (2012), Rapoport, Daniel, and Seale (2008), and especially Danz, Fehr, and Kübler (2012).

24. In the accumulating balances version of the experiment, the same person is above a given player this round as in the previous round 56% of the time, compared to 15% in the rerandomized one-shot version of the game.

the player below results in a drop in rank. Instructions and a typical screen shot from the game are found in the Online Appendix. As the choice between the person directly above and below is not well defined for the first- and last-place players, we have the first-place player choose between the second- and third-place player, and the last-place player between the fourth- and fifth-place player. The choice sets are summarized in Online Appendix Table 6. Players are clearly instructed that the additional \$2 comes from a separate account and not from the player herself.

After players make their decisions, one player is randomly chosen and his choice determines the final payoffs of that round. As such, players should make their decisions as if they alone will determine the final distribution of the round. To avoid any reciprocity effects, players do not know which player is chosen or the outcome of the round. After the end of each round, players are rerandomized across the same \$1, \$2, . . . , \$6 distribution and the game repeats. They are paid their final balances for one randomly chosen round.

IV.B. Separating LPA from Alternative Models

Inequality aversion as in Fehr and Schmidt (1999) predicts that all players give to the lower-ranked player. In fact, we designed the experiment so that the net effect of giving to the lower-ranked person with respect to the standard Fehr-Schmidt inequality terms is constant for ranks one through five.²⁵ As such, although inequality-aversion makes the prediction that players should generally give to the lower-ranked player, it predicts that this tendency should be no different for those close to the bottom of the distribution. LPA, by contrast, predicts that players near the bottom will give to the lower-ranked person less often.

There is strong empirical support for subjects generally favoring a fellow subject with less money. In the experiments of Engelmann and Strobel (2004), giving to the lower-ranked player involves substantially lowering total surplus, but subjects do so regardless about half the time. As Tricomi et al. (2010) show, in both subjective ratings and functional magnetic resonance imaging data, the poorer member in a two-player game evaluates transfers to the richer member more negatively than the richer

25. The note under Online Appendix Table 6 shows the simple arithmetic behind this claim.

person evaluates transfers to the poorer person. LPA thus requires that those at the bottom of the distribution overcome the psychological cost typically associated with giving money to someone richer.

Because we hold total surplus (the \$2 must go to someone) and own income constant, we are able to separate LPA from several other models of distributional preferences. First, many papers have posited that utility is a positive function of $\frac{y_i}{\bar{y}}$, but an individual's decision in our experiment cannot affect either the numerator (she cannot keep the money herself) or the denominator (total surplus is fixed so the average among all players, \bar{y} , is also fixed). Similarly, in Bolton and Ockenfels (2000), utility is based on own income and one's share of total surplus, neither of which is affected by the player's decision to give \$2 to the person above or below. In Charness and Rabin (2002), utility is a weighted sum of own income, total surplus, and the income of the poorest person. Their model in fact predicts that the last-place player will be the most likely to give to the lower-ranked player, as only he can improve the minimum-income level of the distribution by giving \$2 to the person below him.

By design, giving to the lower-ranked player in their choice set causes all players except the first and last to drop one rank in the distribution. We thus predict that first- and last-place players will have the highest rates of giving to the lower-ranked player, as they do not face an equality-rank trade-off. Among those facing such a trade-off (ranks two through five), LPA predicts that dropping in rank would have the largest psychic cost for those close to last place themselves and thus that individuals will be the least likely to give to the lower-ranked player when they themselves are in second-to-last place. Adding LPA to our knowledge of individuals' behavior in simpler redistribution experiments, we predict a strong overall tendency to give to the lower-ranked player, but a substantial reduction in this tendency for those in the bottom of the distribution, particularly for the person in second-to-last place.

IV.C. Initial Results

Our first version of the redistribution experiment grouped players into groups of six, to follow the lottery experiments. Figure IV shows how the probability a player gives the additional \$2 to the lower-ranked player in his choice set varies by rank.

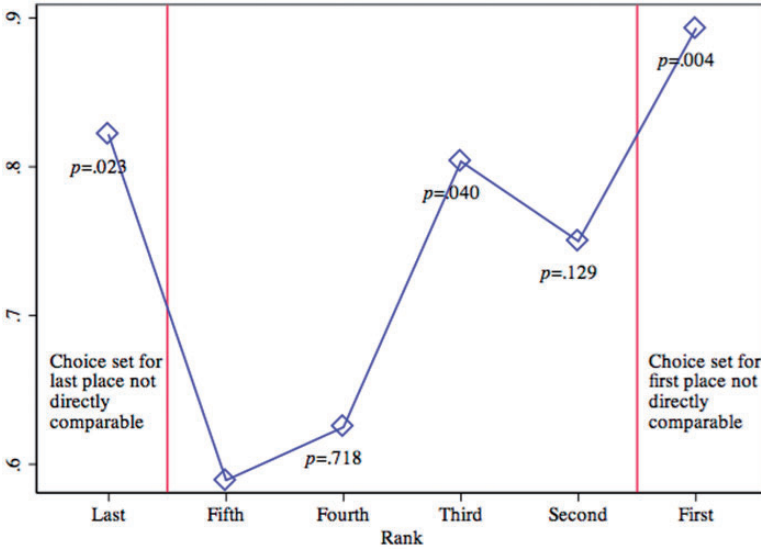


FIGURE IV

Probability of Choosing to Give \$2 to the Lower-Ranked Player in Their Choice Set

Based on seven six-player games of eight rounds each, giving a total of 336 observations. Each player except the first- and last-place player were given the choice between giving an extra \$2 to the person directly above or below them in the distribution. The first-place player decided between the second- and third-place player, while the last-place player decided between the fourth- and fifth-place player. See Section IV.A for details. All *p*-values are based on the probit regression: $gave\ to\ lower\ rank_i = \sum_{k \neq 5} \beta^k rank_i^k + \epsilon_i$, where $rank_i^k$ is an indicator variable for player *i* having rank *k*, standard errors are clustered by player, and no other controls are included. As there are six ranks in this version of the redistribution experiment, the excluded category is fifth (second-to-last) place. The *p*-values in the figure refer to the estimated fixed effect of being in rank *k* relative to the excluded category of fifth place. The *y*-axis values are the probit coefficients (as changes in probability) plus the mean of the excluded category (to normalize).

Overall, players choose to give to the lower-ranked player in their choice set 75% of the time, consistent with inequality aversion. This probability varies from over 80% in the top half of the distribution, to less than 60% for the second-to-last-place player. Players are the least likely to give to the last-place player when they are in second-to-last place and this difference is pairwise significant for the first-, third-, and last-place players, and marginally significant (*p* = .120) for the second-place player.

Those third-from-last (for whom giving to the lower-ranked player leads to a demotion to second-to-last) are nearly equally likely to deny the \$2 to the lower-ranked player, though the difference between the second- and third-from-last subjects is slightly more pronounced when the first two rounds are dropped (Online Appendix Figure 4).

The first- and last-place players are the most likely to give to the lower-ranked player in their choice set, consistent with their not facing an equality-rank trade-off. Interestingly, the last-place player is relatively more likely to give to the higher-ranked player, perhaps because giving to the second-to-last-place player means he is more isolated in last place.

Table II presents probit regression results reported as changes in probability. In all cases, round and game fixed effects and separate dummy variables for the first- and last-place players are included, since these two players do not have parallel choice sets to those of other ranks. Column (1) shows that the second-to-last-place player is significantly less likely to give to the lower-ranked player relative to other players, and column (2) shows that the the same pattern holds if the second- and third-from-last players are grouped into one category.

A key challenge in separating any LPA effect from competing hypotheses is that with only six ranks we have limited degrees of freedom. This problem is aggravated in the current experiment relative to the the lottery experiment because only ranks two through five have comparable choice sets, whereas in the lottery game we could compare ranks one through six. Being able to compare only four ranks makes it impossible to separate, say, a story in which individuals dislike being near last place versus one in which they want to be above the median. For this reason, we rerun the experiment with eight players.

IV.D. Results from the Eight-Player Game

Beyond the number of players, the game is exactly parallel to the six-player game described in Section IV.A. Players ($N = 72$, divided into nine eight-player games) in ranks two through seven must decide between giving \$2 to the person directly above them or below them, and the first-place player decides between the second- and third-place players whereas the last-place player decides between the sixth- and seventh-place players.

TABLE II
PROBIT REGRESSIONS OF THE PROPENSITY TO GIVE \$2 TO THE LOWER-RANKED PLAYER

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
	Six players			Eight players										
Second from last	-0.116* [0.064]		-0.071* [0.040]	-0.097** [0.046]	-0.068** [0.030]	-0.060 [0.044]	-0.095** [0.048]	-0.081* [0.049]	-0.072 [0.053]	-0.104* [0.058]	-0.150* [0.085]	-0.090** [0.035]		
Second or third from last		-0.147** [0.061]											-0.097** [0.030]	-0.154** [0.067]
Below median						-0.017 [0.036]	-0.005 [0.042]	0.018 [0.053]						
Rank, scaled 0 to 1									0.002 [0.088]	0.016 [0.101]	0.150 [0.178]			0.135 [0.139]
Mean, dept. var.	0.747 No	0.747 No	0.802 No	0.784 Yes	0.802 All	0.802 All	0.784 Yes	0.802 No	0.802 No	0.784 Yes	0.784 Yes	0.784 No	0.784 No	0.784 No
Ex. early rds?														
Log likelihood	-167.6	-165.5	-306.8	-252.4	-306.2	-306.7	-252.3	-306.2	-306.8	-252.3	-252.3	-478.9	-476.6	-475.9
Observations	336	336	648	504	648	648	504	648	648	504	504	984	984	984

Notes. All regressions are estimated via probit, coefficients are reported as marginal changes in probability, round and game fixed effects are included, and standard errors are clustered by individual. The first two columns are based on seven six-player games of eight rounds each, giving a total of 336 observations, and the next nine columns are based on nine eight-player games of nine rounds each, giving a total of 648 observations, and the final three columns pool all observations. The dependent variable is an indicator variable for whether the individual chose to give \$2 to the lower ranked of the two players in his choice set. Specifications that “exclude early rounds” drop the first two rounds. See Section IV for further details on the experiment. * $p < .10$, ** $p < .05$, *** $p < .01$.

Figure V presents the basic results from the eight-player game. As before, the second-to-last-place player is the least likely to give to the player below him, and this difference is often pair-wise significant from other ranks. Also as before, the third-to-last-place player appears similar to the last-place player. Importantly, however, the player just below the median (rank = 5) shows no such tendency, and the pairwise difference with the second-to-last-place player is statistically significant. Put differently, comparing the six- and eight-player games suggests that there is nothing particularly salient about being, say, in fourth or fifth place; instead behavior appears to depend on how close one is to last place: the fourth- and fifth-place players in the six-player game show strong evidence of LPA, while the fourth- and fifth-place players in the eight-player game do not.

Columns (3) through (11) of Table II present results from the eight-player game. Consistent with the figure, in column (3) the second-to-last-place player is significantly less likely to give to the lower-ranked player than are other players (again, the first- and last-place players always have their own fixed effect, so their generally higher tendency to give to the lower-ranked player does not contribute to the coefficient), and this effect increases when early rounds are excluded (column (4)). In column (5) we gain precision (the standard error falls by one-fourth) by including those in third-to-last place as being affected by last-place aversion: if they give \$2 to the lower-ranked player, they would fall into second-to-last place.

Column (6) explores whether LPA can be instead explained by individuals simply wanting to be above the median. Although the p -value of the second-to-last-place term is not quite significant ($p = .112$), it becomes significant when the first two rounds are excluded (column (7)) or if we also include the third-to-last-place player to gain precision (column (8)).

Controlling for rank actually increases the coefficient on the second-to-last term (column (9)), though, as with the lottery experiments, the high multi-collinearity between rank and the variable of interest significantly increases the standard errors. When we exclude the first two rounds (column (10)) or pool the second- and third-from-last players (col. (11)), the effect regains its significance. In fact, the coefficient on rank is “wrong-signed” in all of our tests (higher rank tends to increase giving to the lower-ranked player, thus cutting in the opposite direction as LPA), and thus adding it always increases the LPA effects. In columns (12)




FIGURE V

Probability of Choosing to Give \$2 to the Lower-Ranked Player in Their Choice Set (Eight-Player Game)

Based on nine eight-player games of nine rounds each, giving a total of 648 observations. Each player except the first- and last-place player were given the choice between giving an extra \$2 to the person directly above or below them in the distribution. The first-place player decided between the second- and third-place player, and the last-place player decided between the sixth- and seventh-place player. See Section IV.D for details. All p -values are based on the probit regression: $gave\ to\ lower\ rank_i = \sum_{k \neq 7}^8 \beta^k rank_i^k + \epsilon_i$, where $rank_i^k$ is an indicator variable for player i having rank k , standard errors are clustered by player, and no other controls are included. As there are eight ranks in this version of the redistribution experiment, the excluded category is seventh (second-to-last) place. The p -values in the figure refer to the estimated fixed effect of being in rank k relative to the excluded category of seventh place. The y -axis values are the probit coefficients (as changes in probability) plus the mean of the excluded category (to normalize).

to (14) we pool the six- and eight-player experiments and show that we can separate a linear rank effect from low-rank aversion more definitively with this larger sample.

Online Appendix Table 7 shows that the results are robust to demographic controls and presents some differential treatment effects. Interestingly, self-identified religious and politically conservative people show stronger LPA effects. Such individuals are significantly undersampled in our experiment relative to the

general population, suggesting a more representative sample might display even larger LPA effects.²⁶

As noted earlier, inequality aversion in the standard two-term Fehr-Schmidt parameterization cannot explain our results, as the decision to give the \$2 to the person above or below has the same net effect on their inequality-aversion terms regardless of rank. We thus experiment with alternative measures of inequality-aversion and social comparison. Whereas Fehr and Schmidt focus on the total income above and below, individuals may instead focus on the average income of those above and below them. Or, individuals may try to maximize their position within the income range (Brown et al. 2008) or their position in the range relative to the last-place person, $\frac{v_i - v_{last}}{Range}$ (Rablen 2008). Alternatively, they may wish to minimize the Gini coefficient of the distribution. Online Appendix 8 shows that LPA is robust to each of these controls, and in addition is robust to controlling for one's rank in the previous round.

IV.E. Discussion

The results from these experiments offer broad support for the hypothesis that players experience disutility from being in the bottom of the distribution. This effect can be separated from individuals' merely wanting to be above the median as well as inequality aversion, surplus maximization, and linear controls for rank. This result in fact contradicts the predictions of maximin models.

Both the six- and eight-player games suggest that players take action to avoid falling not just to the very bottom rank but to the second-lowest rank as well. Two possible explanations seem likely. First, players may have a similar distaste for being near last place in a distribution as they do for being in last place itself. In both experiments, this heightened concern over rank appears to diminish once players are safely near the middle of the distribution. Alternatively, they may care only about avoiding last place, but may have mistakenly played the game as strategic when, because only one randomly chosen player's decision is

26. For example, for the GSS question asking respondents to place themselves on a 7-point conservative-to-liberal scale, the average is 4.11, compared to 5.3 (5.4) in the six- (eight-) player distribution games (see summary statistics in Online Appendix Table 1). Similarly, 18% of people in the GSS describe themselves as "very religious," compared to 4% of our experimental sample.

implemented, it is actually nonstrategic. To facilitate data collection, we had players choose “as if” they were the dictator, but recent work has found that “role uncertainty” can have modest effects on players’ decisions, even when it should not in principle.²⁷ Especially in early rounds, when the pattern of choosing between the person above and below you is less apparent, the third-from-last player may have assumed that everyone else would give money to the last-place player, and thus (incorrectly) inferred that by allowing the second-to-last-place player to leapfrog him, he would run the risk of falling to last place himself. In any case, as we predicted, players appear less willing to sacrifice rank when they are already near the bottom of the distribution.

V. LAST-PLACE AVERSION AND SUPPORT FOR MINIMUM WAGE INCREASES

In choosing a real-world policy to test the predictions of LPA, we begin with the minimum wage. First, the minimum wage defines the “last-place” wage that can legally be paid in most labor markets, so it allows us to define “last place” more easily than in the context of other policies. Second, although the worst-off workers are not always those being paid the minimum wage (e.g., middle-class teenagers might take minimum-wage jobs during the summer), previous research has shown that policies that more explicitly target the poor (such as Temporary Assistance for Needy Families) could have potentially confounding racial associations (though we briefly examine welfare support in the next section).²⁸

We emphasize up front that actual policies simply do not have the same power to reject alternative distributional preferences that our redistribution experiment in the previous section does. No policy asks individuals to choose between helping those directly above or below them, and most policies that respondents would recognize as redistributive generally involve helping those

27. As Iriberry and Rey-Biel (2011) note, role uncertainty has been found to encourage “strategic thinking” in games, consistent with third-from-last-place players thinking they may need to defend against others’ generosity toward the last-place player. Engelmann and Strobel (2007) also find differences in dictator games with and without role uncertainty. We thank Doug Bernheim for alerting us to this possibility.

28. See Gilens (1996).

at the bottom of the distribution. As such, we view the evidence in this and the following section as testing whether preferences are consistent with LPA, but we are aware that the results cannot eliminate all alternative theories.

V.A. Predicting Who Would Support a Minimum Wage Increase

A minimum wage increase is a transfer to some low-wage workers from—depending on market characteristics—other low-wage workers who now face greater job rationing, employers with monopsony power in the labor market, or consumers who now pay higher prices.

Assuming low-wage workers are not concerned with adverse employment effects—a hypothesis we directly test in the empirical work—they should generally exhibit the greatest support for an increase relative to other workers. First, they themselves might see a raise, depending on the difference between their current wage and the proposed new minimum and the strength of spillover effects to workers just above the proposed new minimum.²⁹ Second, even for those who would not be directly affected, the policy could act as wage insurance and should increase their reservation wage. Finally, if low-wage workers are relatively substitutable, then those making just above the current minimum should welcome a minimum-wage increase because employers would then have less opportunity to replace them with lower-wage workers.

LPA, in contrast, predicts that individuals making just above the current minimum would have limited enthusiasm for seeing it increased. The minimum wage essentially defines the last-place wage a worker in most labor markets can legally be paid. A worker making just above the current minimum might see a wage increase from the policy, but could now herself be tied with many other workers for last place.

V.B. Minimum Wage Survey Data

Questions regarding the minimum wage have often appeared in opinion surveys, but to the best of our knowledge none have also asked respondents to report their own wages (as opposed to

29. The strength of spillover effects has been debated in the literature. Lee (1999) found evidence consistent with large spillovers, whereas Autor, Manning, and Smith (2010) more recently found far more modest spillovers and in fact showed that they could be due entirely to measurement error.

household income). We thus designed our own survey, which was in the field twice (November and December 2010).³⁰ Subjects were randomly selected from a nationwide pool and invited to complete the online survey in exchange for \$5. Enrollment in the study was limited to employed individuals between the ages of 23 and 64, so as to target prime-age workers. We also over-sampled low-wage and hourly workers.

The survey stated the current federal minimum wage (\$7.25) and then asked respondents whether it should be increased, decreased, or left unchanged. As only 2% wished it to be decreased, our main outcome variable is an indicator for wanting it increased, as opposed to decreased or maintained at the current level.

One version of the survey only sampled hourly workers, whom we asked: “What is your current hourly wage? If you have more than one job, please enter the wage for your main job.” For the other survey, which sampled both hourly and salaried workers, we follow the Congressional Budget Office in their generation of the U.S. wage distribution by asking respondents to divide their paycheck by their usual hours in a pay period to calculate an effective hourly wage: “Even if you are not actually paid by the hour, please calculate your estimated hourly wage. You can do this by dividing your paycheck by how many hours you typically work in a pay period.”³¹ Although we did not ask this sample whether they were hourly workers, given their effective hourly wages the large majority are likely hourly.³²

We make the following sampling restrictions in generating our regression sample. First, we drop the 74 people who completed the survey in less than two minutes (even though we wrote the survey, it took us more than three minutes to complete). We also drop from this sample 12 individuals who report being unemployed but somehow slipped through the survey’s filter.

30. The survey was administered by C&T Marketing Group, <http://www.ctmarketinggroup.com>.

31. See note 1 of Congressional Budget Office, “Changes in the Distribution of Workers’ Hourly Wages between 1979 and 2009,” February 2011, available at <http://www.cbo.gov/sites/default/files/cbofiles/ftpdocs/120xx/doc12051/02-16-wage-dispersion.pdf>.

32. We use the 2011 Current Population Survey to calculate that of those adults making an effective wage of less than \$20 an hour (very similar to our Internet sample), 72% were hourly employees. Unlike our survey, the CPS directly asks workers if they are hourly.

We also drop 63 observations with missing or unusable wage data (e.g., “Depends”). These exclusions leave a regression sample of 489 observations, with a median wage of \$13.80. Online Appendix Table 9 displays summary statistics from the final regression sample.³³

V.C. Results

Figure VI shows how support for increasing the minimum wage varies across wage groups. We break up the distribution into \$1 bins above the current minimum wage of \$7.25, given past work showing that individuals tend to think in \$1 increments.³⁴ As in past surveys, increasing the minimum wage is a popular policy—roughly 80% of our sample appears to support the idea. The striking exception, however, is the relative lack of support among those making just above the current minimum (between \$7.26 and \$8.25). They are, in fact, the group least likely to support it, and the difference between them and other groups in the figure is often statistically significant. With the exception of this group, support for increasing the minimum wage appears to decrease roughly linearly with individuals’ wages.³⁵

Table III presents probit regression results. Column (1) includes a dummy for being just above the current minimum wage (i.e., making more than \$7.25 but no more than \$8.25 an hour) and a linear wage control. We top-code the linear wage control at the 90th percentile because of some unrealistically high reported wage levels that we suspect were caused by missed decimal points.³⁶ Without any other controls, the “just above” coefficient

33. Relative to the 2011 American Community Survey (ACS), our survey oversamples women and the college-educated. In Online Appendix Table 12 we replicate our main results reweighting observations to match the ACS along these two dimensions, producing very similar results.

34. See Basu (1997) and citations therein.

35. As noted already, we know of no survey that allows us to examine minimum wage support by respondent’s own wage, but we did pool recent Pew surveys on the minimum wage to examine support for the minimum wage by income (see Online Appendix Figure 6). Though the analogy to our Figure VI is inexact given the difficulty in relating household income and own wage, the shape is strikingly similar in that low-income groups have relatively soft support for increasing the minimum wage.

36. For example, we strongly suspect that \$1,350 an hour is actually \$13.50 an hour. Our results are robust to dropping top-coded observations or logging (uncapped) wage levels (Online Appendix Tables 10 and 11, respectively).

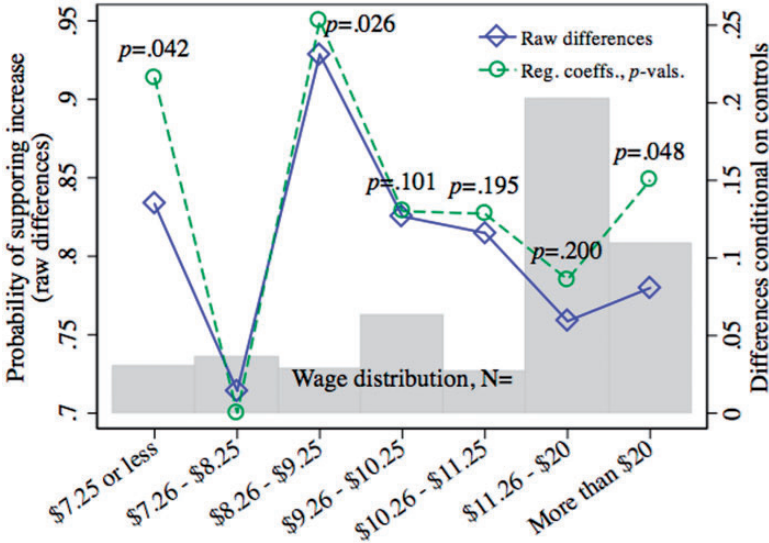


FIGURE VI

Support for Increasing the Minimum Wage from \$7.25, by Wage Rate

Based on the authors' online survey of employed individuals ages 23 to 64. See Section V.B for details. The first series displays the share of each wage group that supports increasing the minimum wage. The second series plots the coefficients (with *p*-values labeled) on the wage-category fixed effects (omitted category $\$7.25 > wage \geq \8.25) from a probit regression that also controls for gender, race, ethnicity, educational level, party affiliation, marital and parental status, approval rating of President Barack Obama, and union status.

is negative but not significant, as shown in column (1). Note that this is a fairly demanding specification because those just above the minimum wage are being compared with individuals with incomes far greater than theirs who are likely different on many important dimensions—when in column (2) we limit the sample to those below the median wage, the effect of being just above the current minimum becomes significant, not surprising given the striking pattern in Figure VI.

In column (3), adding basic demographic controls substantially increases the effect of being just above the minimum wage. This result is not surprising—the types of workers who normally support a minimum wage increase (women, minorities, young workers) are overrepresented among those making just above \$7.25. Adding controls for census division, the state-level minimum wage and an indicator for whether it is above \$7.25

TABLE III
PROBIT REGRESSIONS OF THE PROPENSITY TO SUPPORT A MINIMUM WAGE INCREASE

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Just above min. wage (\$7.26 - 8.25)	-0.0921 [0.0703]	-0.164** [0.0785]	-0.135* [0.0713]	-0.149** [0.0704]	-0.151** [0.0705]	-0.138** [0.0663]	-0.142** [0.0642]
Hourly wage	-0.00217 [0.00183]	-0.0337*** [0.0141]	-0.00114 [0.00189]	-0.00270 [0.00190]	-0.00161 [0.00214]	-0.00241 [0.00204]	-0.00209 [0.00203]
Male			-0.0687* [0.0402]	-0.0558 [0.0400]	-0.0680* [0.0406]	-0.0541 [0.0390]	-0.0496 [0.0379]
Black			0.306** [0.126]	0.310** [0.129]	0.288** [0.132]	0.125 [0.128]	0.136 [0.126]
Hispanic			-0.105 [0.127]	-0.0902 [0.128]	-0.102 [0.118]	-0.0887 [0.111]	-0.0709 [0.111]
Age div. by 100			-0.166 [0.176]	-0.133 [0.176]	-0.0501 [0.190]	0.0574 [0.185]	0.0454 [0.181]
Native born			0.0448 [0.0948]	0.0782 [0.0961]	0.0953 [0.0965]	0.0995 [0.0927]	0.109 [0.0882]
Min. wage threatens job							-0.0433*** [0.00947]
Mean, dept. var	0.785	0.787	0.786	0.784	0.784	0.784	0.784
Sample	All	Low-wage	All	All	All	All	All
Geogr. controls	No	No	No	Yes	Yes	Yes	Yes
Backgr. controls	No	No	No	No	Yes	Yes	Yes
Polit. controls	No	No	No	No	No	Yes	Yes
Log likelihood	-253.1	-122.6	-244.1	-233.0	-229.9	-208.8	-199.1
Observations	489	244	486	481	481	481	481

Notes: All data are from the minimum wage survey (see Section V B for further detail) and all regressions are probit regressions (coefficients reported as marginal changes in probability) for whether a respondent answered that the minimum wage should be increased. The wage control is top-coded at \$46, the 90th percentile, because of accuracy concerns for very high wages (see note 7). In column (2), individuals with wages above the median of \$13.80 are excluded. In column (4), "geographic controls" include fixed effects for the eight census divisions, the level of the state minimum wage, and an indicator variable for whether the state minimum is above the federal minimum. In column (5), "background controls" include marital status and indicator variables for no high school, some high school, high school degree, some college, two-year college degree, four-year college degree, master's degree, doctoral degree, professional degrees. In column (6), "political controls" include fixed effects for major party affiliation; a one-to-seven approval rating of President Barack Obama; and union status. Column (7) controls for respondents' answer to the following question: "Do you worry that if the minimum wage is set too high, it might make employers reduce hiring and possibly cause you to lose your job?" where 1 indicates "not at all worried" and 7 indicates "very worried." * $p < .10$, ** $p < .05$, *** $p < .01$.

marginally increases the coefficient of interest (column (4)). Similarly, controlling in column (5) for education and marital status also marginally increases the magnitude of the effect, as does controlling for party affiliation, union status, and approval rating of President Barack Obama (column (6)).

An important concern is that our results may be driven by low-wage workers' fear that a minimum wage increase will cause disemployment. For this reason, we also asked participants: "Do you worry that if the minimum wage is set too high, it might make employers reduce hiring and possibly cause you to lose your job?" We control for the answer to this question in column (7). Not surprisingly, those who say a minimum wage increase might threaten their job are far less likely to support an increase, but controlling for this variable does not affect the coefficient on the variable of interest.

Although we tend to think that making within \$1 of the minimum wage is the most natural definition of being just above the minimum wage, Online Appendix Table 13 shows that the main result holds if we change the \$8.25 maximum value by \$0.25 in either direction. Online Appendix Table 14 shows the results are robust to using an ordered probit model across the three choices of decreasing, maintaining, or increasing the minimum wage.

While inequality aversion would predict that everyone would support a minimum wage increase and thus cannot explain any deviations, just as we did in the redistribution experiment, we explore whether individuals are averse to seeing the average income of those below them increase.³⁷ In Online Appendix Table 15 we use the Current Population Survey to simulate how different minimum-wage increases would affect the average wage of those below each of our respondents. Controlling for this measure does not appreciably change the main result.

The relatively tepid support among low-income workers for such a transfer is consistent with LPA, as those who are marginally better off seek to retain their ability to distinguish themselves from those in last place. Moving from the laboratory—where reference groups are fixed and highly salient—to the field—where

37. Regarding Fehr-Schmidt-type inequality aversion, for any given worker, an increase in the minimum wage can only decrease the advantageous inequality below her (which the model assumes is desirable) and because it cannot conceivably lead to workers leap-frogging over her, should have no effect on disadvantageous inequality.

individuals can be members of many peer groups—could have diminished LPA. By contrast, the minimum wage results suggest that the income or wage distribution is salient to individuals in the bottom of the distribution, resulting in behavior consistent with the behavior observed in the laboratory experiments.

VI. EVIDENCE OF LPA FROM THE GENERAL SOCIAL SURVEY

The minimum wage has the advantage of creating a clear last-place group but the disadvantage of being a relatively narrow policy. Our final analysis explores general redistributive preferences as a function of income in the General Social Survey (GSS).

VI.A. *Data and Empirical Framework*

The GSS is a repeated cross-sectional survey of around 1,500 respondents a year, conducted every one to two years. The GSS has asked many questions related to redistributive preferences, but many are only asked for a single year. We follow the literature and focus on the GSS question that asks individuals to place themselves on the following scale: “Some people think that the government in Washington should do everything to improve the standard of living of all poor Americans (they are at point 1 on this card). Other people think it is not the government’s responsibility, and that each person should take care of himself (they are at point 5).”³⁸ This question is asked most years from 1975 to 2010. We subtract this variable from 6 so that it is increasing in support for redistribution.

We use the household income variable in the GSS and adjust it to 2011 dollars.³⁹ To account for the fact that one’s taxes and transfers are related to household income adjusted by household size, we follow recent Organisation for Economic Co-operation and Development (OECD) publications and divide household

38. For example, see Alesina and Giuliano (2011).

39. About 10% of respondents do not answer the question, far less than in the European Social Survey or World Values Survey, and refusal does not appear related to redistributive preferences. When we regress our redistribution measure on a dummy for not answering the income question, the point estimate suggests that those who do not answer are slightly more supportive of redistribution, but the p -value is .460.

income by the square root of the number of household members. For ease of exposition we refer to this measure as “income.”⁴⁰

In the redistribution experiments in Section IV, we saw that those in second-to-last place were hesitant to help those below them. We create a rough GSS analogue to this group: those with income above the bottom quintile but below the median (i.e., the sixth through eighth deciles). The analysis below explores whether this group supports redistribution less than one would otherwise predict from their place in the income distribution.

VI.B. Results

Figure VII shows how support for redistribution varies across income deciles. Not surprisingly, there is a negative effect of income decile on redistributive support. However, the deviation from trend among those in the seventh and eighth deciles is quite striking and in general yields a negative but convex relationship. For example, the views of individuals in the seventh decile are as close to those of the first (richest) decile as they are to the tenth (poorest). The results are even more pronounced with other proxies of redistributive support. With respect to increasing the generosity of welfare, Online Appendix Figure 7 shows that those in the eighth decile have closer views to the richest decile than they do to the poorest. With respect to voting for the more redistributive party (the Democrats) in presidential elections, Online Appendix Figure 8 shows those in the seventh decile have substantially closer preferences to those in the richest decile than they do to those in the poorest decile.

It is important to emphasize that this convex relationship, while consistent with LPA, is at odds with classic models of redistributive preferences (e.g., Meltzer and Richard 1981). In these models, individual i gains $\tau\bar{y}$ (where \bar{y} is the mean of the income distribution and τ the tax rate) in transfers while paying τy_i in taxes, so support for redistribution is proportional to $\bar{y} - y_i$. Because the distribution of income y tends to be right-skewed, the empirical relationship between income decile and support for redistribution would be concave. Online Appendix Figure 8 graphs $\bar{y} - y_i$ by income decile in the GSS—in sharp contrast to Figure VII, those in the bottom half of the distribution should

40. See <http://www.oecd.org/statistics/302013041e.pdf> (accessed December 6, 2013).




FIGURE VII

Support for Redistribution over the Income Distribution

Based on data from the General Social Survey. This figure plots coefficients from a regression of $Redistribution_{it} = \sum_n \beta_n \mathbb{1}_i^n + \delta_t + \epsilon_{it}$, where $\mathbb{1}_i^n$ is an indicator variable for being in the n th income decile (the top decile is the omitted category and thus has a coefficient of 0) and δ_t are year fixed effects. Deciles are based on adjusted household income (real household income divided by the square root of household size). *Redistribution* is based on the GSS question *help poor*, which asks individuals to put themselves on a five-point scale (with the extremes being that “government should do everything it can to help the poor” versus “people should fend for themselves”).

have very similar views on redistribution under the classic formulation.

Table IV presents regression analysis. Column (1) includes no controls except the LPA dummy (being in the sixth through eighth deciles) and income decile. Income decile has a negative effect on support for redistribution, but, as predicted, there is a significant deviation from trend for those in the LPA group. In fact, the differential decrease in support for redistribution associated with being in the LPA group (−0.0958) is larger than the decrease associated with moving up an income decile (−0.0870).

This effect is essentially identical when year and region fixed effects are added in column (2). Column (3) shows that the result holds when in addition to controlling for income decile, we also

TABLE IV
SUPPORT FOR REDISTRIBUTION BY INCOME DECILE

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Support redistribution			Welfare	Vote Dem		Redistr.
LPA group (6th through 8th deciles)	-0.0958*** [0.0173]	-0.0947*** [0.0172]	-0.0950*** [0.0172]	-0.0643*** [0.00540]	-0.0347*** [0.00670]	-0.119*** [0.0446]	-0.109** [0.0445]
Income decile	-0.0870*** [0.00277]	-0.0896*** [0.00278]	-0.0887*** [0.00454]	-0.0277*** [0.000876]	-0.0280*** [0.00110]	-0.0981*** [0.00715]	-0.0889*** [0.00741]
Income ÷ 10,000			-0.000554 [0.00216]				
Ever on govt. assistance							0.201*** [0.0434]
Mean, dept. var	3.117	3.117	3.117	0.204	0.467	3.209	3.209
Region FE?	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	23,627	23,627	23,627	27,975	31,708	4,066	4,066

Notes. All regressions include year fixed effects. The redistribution question asks individuals to place themselves on a five-point scale on whether government should help the poor or individuals should fend for themselves. "Increase welfare" is a binary variable indicating that welfare benefits should increase. "Vote Dem" is a binary variable indicating that you voted for the Democratic candidate in the most recent presidential election. To maximize sample size, "Ever receive gov. assistance" is based on two GSS questions. For most observations, it is defined by the GSS variable GOVAID, based on the question: "Did you ever—because of sickness, unemployment, or any other reason—receive anything like welfare, unemployment insurance, or other aid from government agencies?" In 1986, the GSS specifically asked if the respondent was ever on welfare, and we combine this variable in our "government assistance" measure. In most cases, variation in the number of observations is driven by variation in the number of respondents who answered the outcome question. However, in the last two columns, we restrict the sample to having nonmissing responses to both the main redistribution question and our "ever receive assistance" variable. * $p < .10$, ** $p < .05$, *** $p < .01$.

control for actual income. Columns (4) and (5) show that the deviation of the LPA group is even more striking (relative to the coefficient on income decile) when support for increasing welfare generosity or voting Democratic serves as the outcome.

An alternative explanation to LPA is that individuals are merely forming preferences based on self-interest. We explore this possibility by accounting for whether the respondent was ever on government assistance. Unfortunately, there is limited overlap ($N = 4,066$) between our main redistribution outcome and this measure, but the LPA coefficient barely moves when our standard regression is estimated on this sample (column (6)) and when the assistance control is added (column (7)).⁴¹

As noted earlier, a serious concern is that racial attitudes confound these results: individuals may not be last-place averse but simply worry that redistribution will differentially help minority groups (who tend to be at the bottom of the distribution), with whom they feel little kinship. Online Appendix Figures 10 and 11 cast doubt on this alternative hypothesis: support for redistribution and welfare varies across income deciles in a very similar manner for minorities as it does for the full sample, though confidence intervals widen because of the smaller sample.⁴² In Online Appendix Table 17, we show that the regression coefficients barely change when the main specification (column (2) of Table IV) is estimated with non-Hispanic whites versus blacks and Hispanics. As Luttmer (2001) has shown, support for welfare among whites falls when they live near blacks, but in fact the LPA effect among whites is strongest in the “whitest” regions of the country, further suggesting that racial resentment is not driving LPA. Finally, we show that whites in the LPA group are not differentially likely to say that the government is doing “too much” for blacks.

Online Appendix Table 18 provides additional checks, showing that the effect holds for the prime working-age subsample and

41. There is more substantial overlap with the welfare and voting outcomes. When a parallel exercise is performed using these outcomes, adding the government assistance variable barely moves the coefficients of interest (see Online Appendix Table 16).

42. The relationship between voting Democratic and income decile among minorities does not take the same shape (see Online Appendix Figure 12), but then almost all minorities vote Democratic, and thus there is little variation in this variable.

after controlling for demographic and background characteristics, as well as opinions on political and cultural questions.

In summary, the GSS results suggest that once respondents are above the bottom income quintile, their support for redistribution grows soft. These results do not appear to be driven by racial resentment or direct self-interest. Respondents below the median but above the bottom quintile often share opinions about aiding low-income individuals that are similar to those in the top part of the distribution.

VII. CONCLUSION

We design experiments that allow us to separately identify an aversion to being in the bottom of a distribution from a large variety of existing models of preferences and social comparison. Individuals randomly placed in the bottom of a distribution are willing to bear risk for the possibility of improving in rank that they are unwilling to bear when placed higher in that distribution. In modified dictator games, between one-half and one-fourth of second-to-last-place players give additional money to the person above rather than below them, whereas such behavior is very rare for individuals in other parts of the distribution. Taken together, the evidence from these experiments suggest that the predictions from standard models of preferences may break down toward the bottom of the distribution.

We then apply the insights from the redistribution experiments to predict respondents' preferences for a particular redistributive policy: increasing the minimum wage. LPA predicts that those making just above the current minimum wage face a trade-off: on one hand, they may receive a raise if the new minimum wage is above their current wage; on the other hand, they may lose their status of having a wage above the last-place group. We conduct a survey and find that support for a minimum wage increase is lowest among those making just above the current minimum. Finally, we show that support in the GSS for redistribution is particularly soft among those who are above the bottom quintile but below the median income—a rough analogue to those in second-to-last place in our redistribution experiments—and in sharp contrast to models where support for redistribution is proportional to the difference between own income and average income.

While we have focused largely on redistributive preferences, LPA may have other applications. Consumer behavior is a natural extension. Past work has noted consumers' tendency to purchase the second cheapest wine on a menu (McFadden 1999), consistent with consumers exhibiting a standard price response but simultaneously avoiding association with the last-place product. In a choice set of three or four, this same tendency would lead consumers to pick a "middle" option—the "compromise effect" in behavioral decision theory (Simonson 1989). Similarly, research on the psychology of queuing has shown that those at the end of a line are the least likely to allow the person behind them to pay to leapfrog them (Oberholzer-Gee 2006), suggesting psychological disutility from being last in line. Interventions targeting the last-place client might improve the performance of service operations in both the private and public sectors.

We examined risk taking and money transfer decisions, but individuals likely cope with the disutility of last place or low rank in other ways. For example, those in last place might work especially hard at a given task to move up in rank, they might instead give up, or they might seek to define themselves along an alternative metric or within a different peer group. Evidence from the field is limited but points in the direction of low rank leading to discouragement in educational settings.⁴³ Boys in the Moving to Opportunity experiment who moved to better neighborhoods and schools, and thus found themselves in the bottom of the classroom academic distribution, actually exhibited an increase in criminal activity relative to the control group (Kling, Ludwig, and Katz 2005)—consistent with low rank in a given domain encouraging individuals to substitute effort to another. Better understanding how individuals cope with low rank could inform interventions that target disadvantaged and at-risk individuals.

COLUMBIA BUSINESS SCHOOL
HARVARD BUSINESS SCHOOL
STANFORD GRADUATE SCHOOL OF BUSINESS
HARVARD BUSINESS SCHOOL

43. Hoxby and Weingarth (2005) use random assignment to classrooms to demonstrate that low-achieving students perform better when they are not the only low-achieving student in the classroom. Similarly, Tran and Zeckhauser (2012) find that randomly informing students of their rank after an exam increases study effort, but only among those with high rank.

SUPPLEMENTARY MATERIAL

An Online Appendix for this article can be found at QJE online (qje.oxfordjournals.org).

REFERENCES

- Alesina, Alberto F., and Paola Giuliano, "Preferences for Redistribution," in *Handbook of Social Economics*, vol. 1A, Matthew O. Jackson, Jess Benhabib, and Alberto Bisin, eds. (Amsterdam: North-Holland, 2011).
- Autor, David, Alan Manning, and Christopher L. Smith, "The Contribution of the Minimum Wage to US Wage Inequality over Three Decades: A Reassessment," Technical report, National Bureau of Economic Research, 2010.
- Basu, Kaushik, "Why Are so Many Goods Priced to End in Nine? And Why This Practice Hurts the Producers," *Economics Letters*, 54 (1997), 41–44.
- Benabou, R., and E. A. Ok, "Social Mobility and the Demand for Redistribution: The POUM Hypothesis," *Quarterly Journal of Economics*, 116 (2001), 447–487.
- Benartzi, Shlomo, and Richard H. Thaler, "Myopic Loss Aversion and the Equity Premium Puzzle," *Quarterly Journal of Economics*, 110 (1995), 73–92.
- Blanchflower, David G., and Andrew J. Oswald, "Well-Being over Time in Britain and the USA," *Journal of Public Economics*, 88 (2004), 1359–1386.
- Bolton, Gary E., and Axel Ockenfels, "ERC: A Theory of Equity, Reciprocity, and Competition," *American Economic Review* (2000), 166–193.
- Boyce, Christopher J., Gordon D. A. Brown, and Simon C. Moore, "Money and Happiness: Rank of Income, not Income, Affects Life Satisfaction," *Psychological Science*, 21 (2010), 471–475.
- Bredemeier, Christian, "Imperfect Information and the Meltzer-Richard Hypothesis," *Public Choice* (2010), 1–16.
- Brown, Gordon D. A., Jonathan Gardner, Andrew J. Oswald, and Jing Qian, "Does Wage Rank Affect Employees' Well-Being?," *Industrial Relations: A Journal of Economy and Society*, 47 (2008), 355–389.
- Camerer, C. F., *Behavioral Game Theory: Experiments in Strategic Interaction* (Princeton, NJ: Princeton University Press, 2003).
- Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong, "A Cognitive Hierarchy Model of Games," *Quarterly Journal of Economics*, 119 (2004), 861–898.
- Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez, "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction," *American Economic Review*, 102 (2012), 2981–3003.
- Carlsson, F., "Design of Stated Preference Surveys: Is There More to Learn from Behavioral Economics?," *Environmental and Resource Economics*, 46, no. 2 (2010), 167–177.
- Charles, Kerwin Kofi, "Erik Hurst, and Nikolai Roussanov, "Conspicuous Consumption and Race," *Quarterly Journal of Economics*, 124 (2009), 425–467.
- Charness, Gary, and Matthew Rabin, "Understanding Social Preferences with Simple Tests," *Quarterly Journal of Economics*, 117 (2002), 817–869.
- Clark, Andrew E., Paul Frijters, and Michael A. Shields, "Relative Income, Happiness, and Utility: An Explanation for the Easterlin Paradox and Other Puzzles," *Journal of Economic Literature*, 46 (2008), 95–144.
- Clark, Andrew E., Westergård-Nielsen Niels, and Nicolai Kristensen, "Economic Satisfaction and Income Rank in Small Neighbourhoods," *Journal of the European Economic Association*, 7 (2009), 519–527.
- Cole, Harold L., George J. Mailath, and Andrew Postlewaite, "Social Norms, Savings Behavior, and Growth," *Journal of Political Economy* (1992), 1092–1125.
- Costa-Gomes, Miguel A., and Georg Weizsäcker, "Stated Beliefs and Play in Normal-Form Games," *Review of Economic Studies*, 75 (2008), 729–762.

- Danz, David N., Dietmar Fehr, and Dorothea Kübler, "Information and Beliefs in a Repeated Normal-Form Game," *Experimental Economics*, 15 (2012), 622–640.
- Dohmen, Thomas, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert Wagner, "Individual Risk Attitudes: New Evidence from a Large, Representative, Experimentally-Validated Survey," Technical report, IZA discussion paper, 2005.
- Duesenberry, J. S., *Income, Saving and the Theory of Consumer Behavior* (Cambridge, MA: Harvard University Press, 1949).
- Durante, Ruben, Louis Putterman, and Joël van der Weele, "Preferences for Redistribution and Perception of Fairness: An Experimental Study," *Journal of the European Economic Association* (forthcoming).
- Eaton, B. Curtis, and Mukesh Eswaran, "The Evolution of Preferences and Competition: A Rationalization of Veblen's Theory of Invidious Comparisons," *Canadian Journal of Economics / Revue canadienne d'économique*, 36 (2003), 832–859.
- Engelmann, D., and M. Strobel, "Inequality Aversion, Efficiency, and Maximin Preferences in Simple Distribution Experiments," *American Economic Review*, 94 (2004), 857–869.
- , "Preferences over Income Distributions Experimental Evidence," *Public Finance Review*, 35 (2007), 285–310.
- Fehr, E., and K. M. Schmidt, "A Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, 114 (1999), 817–868.
- Frank, Robert H., "The Demand for Unobservable and Other Nonpositional Goods," *American Economic Review*, 75 (1985), 101–116.
- Gilens, M., "Race Coding" and White Opposition to Welfare," *American Political Science Review* (1996), 593–604.
- Gneezy, U., and J. Potters, "An Experiment on Risk Taking and Evaluation Periods," *Quarterly Journal of Economics*, 112 (1997), 631–645.
- Hamermesh, Daniel S., "Interdependence in the Labour Market," *Economica*, 42, no. 168 (1975), 420–429.
- Harbaugh, William T., Kate Krause, and Lise Vesterlund, "Risk Attitudes of Children and Adults: Choices over Small and Large Probability Gains and Losses," *Experimental Economics*, 5 (2002), 53–84.
- Harrison, Glenn W., John A. List, and Charles Towe, "Naturally Occurring Preferences and Exogenous Laboratory Experiments: A Case Study of Risk Aversion," *Econometrica*, 75 (2007), 433–458.
- Heinemann, Frank, "Measuring Risk Aversion and the Wealth Effect," *Research in Experimental Economics*, 12 (2008), 293–313.
- Holt, Charles A., and Susan K. Laury, "Risk Aversion and Incentive Effects," *American Economic Review*, 92 (2002), 1644–1655.
- Hoxby, C. M., and G. Weingarth, "Taking Race out of the Equation: School Reassignment and the Structure of Peer Effects," Unpublished manuscript.
- Hyndman, Kyle, Erkut Y. Ozbay, Andrew Schotter, and Wolf Zeev Ehrblatt, "Convergence: An Experimental Study of Teaching and Learning in Repeated Games," *Journal of the European Economic Association*, 10 (2012), 573–604.
- Iriberrri, Nagore, and Pedro Rey-Biel, "The Role of Role Uncertainty in Modified Dictator Games," *Experimental Economics*, 14, no. 2 (2011), 160–180.
- Kling, J. R., J. Ludwig, and L. F. Katz, "Neighborhood Effects on Crime for Female and Male Youth: Evidence from a Randomized Housing Voucher Experiment," *Quarterly Journal of Economics*, 120 (2005), 87–130.
- Lee, David S., "Wage Inequality in the United States during the 1980s: Rising Dispersion or Falling Minimum Wage?," *Quarterly Journal of Economics*, 114 (1999), 977–1023.
- Lee, Woojin, and John E. Roemer, "Racism and Redistribution in the United States: A Solution to the Problem of American Exceptionalism," *Journal of Public Economics*, 90 (2006), 1027–1052.
- Levy, Haim, "Absolute and Relative Risk Aversion: An Experimental Study," *Journal of Risk and Uncertainty*, 8 (1994), 289–307.
- Luttmer, E. F. P., "Group Loyalty and the Taste for Redistribution," *Journal of Political Economy*, 109 (2001), 500–528.

- , “Neighbors as Negatives: Relative Earnings and Well-Being,” *Quarterly Journal of Economics*, 120 (2005), 963–1002.
- McFadden, D., “Rationality for Economists?,” *Journal of Risk and Uncertainty*, 19 (1999), 73–105.
- Meltzer, A. H., and S. F. Richard, “A Rational Theory of the Size of Government,” *Journal of Political Economy* (1981), 914–927.
- Moore, Don A., John M. Oesch, and Charlene Zietsma, “What Competition? Myopic Self-Focus in Market-Entry Decisions,” *Organization Science*, 18 (2007), 440–454.
- Nagel, Rosemarie, “Unraveling in Guessing Games: An Experimental Study,” *American Economic Review*, 85 (1995), 1313–1326.
- Oberholzer-Gee, Felix, “A Market for Time Fairness and Efficiency in Waiting Lines,” *Kyklos*, 59 (2006), 427–440.
- Ochs, Jack, “Games with Unique, Mixed Strategy Equilibria: An Experimental Study,” *Games and Economic Behavior*, 10 (1995), 202–217.
- Rabin, Matthew, “Risk Aversion and Expected-Utility Theory: A Calibration Theorem,” *Econometrica*, 68 (2000), 1281–1292.
- Rablen, Matthew D., “Relativity, Rank and the Utility of Income,” *Economic Journal*, 118, no. 528 (2008), 801–821.
- Radzevick, Joseph R., and Don A. Moore, “Myopic Biases in Competitions,” *Organizational Behavior and Human Decision Processes*, 107, no. 2 (2008), 206–218.
- Raleigh, Michael J., Michael T. McGuire, Gary L. Brammer, and Arthur Yuwiler, “Social and Environmental Influences on Blood Serotonin Concentrations in Monkeys,” *Archives of General Psychiatry*, 41 (1984), 405.
- Rapoport, A., and P. S. Dale, “The ‘End’ and ‘Start’ Effects in Iterated Prisoner’s Dilemma,” *Journal of Conflict Resolution*, 10, no. 3 (1966), 363.
- Rapoport, Amnon, Terry E. Daniel, and Darryl A. Seale, “Asymmetric Two-Person Bargaining under Incomplete Information: Strategic Play and Adaptive Learning,” *Handbook of Experimental Economics Results*, 1 (2008), 560–571.
- Roemer, John E., “Why the Poor Do Not Expropriate the Rich: An Old Argument in New Garb,” *Journal of Public Economics*, 70 (1998), 399–424.
- Simonson, I., “Choice Based on Reasons: The Case of Attraction and Compromise Effects,” *Journal of Consumer Research* (1989), 158–174.
- Stahl, Dale O., and Paul W. Wilson, “On Players Models of Other Players: Theory and Experimental Evidence,” *Games and Economic Behavior*, 10 (1995), 218–254.
- Tran, Anh, and Richard Zeckhauser, “Rank as an Inherent Incentive: Evidence from a Field Experiment,” *Journal of Public Economics*, 96 (2012), 645–650.
- Tricomi, E., A. Rangel, C. F. Camerer, and J. P. O’Doherty, “Neural Evidence for Inequality-Averse Social Preferences,” *Nature*, 463, no. 7284 (2010), 1089–1091.
- Veblen, T., *The Theory of the Leisure Class* (New York: Macmillan, 1899).
- Zizzo, Daniel John, “Between Utility and Cognition: The Neurobiology of Relative Position,” *Journal of Economic Behavior & Organization*, 48 (2002), 71–91.

This page intentionally left blank