

What do Americans know about inequality? It depends on how you ask them

Kimmo Eriksson*[†]

Brent Simpson[‡]

Abstract

A recent survey of inequality (Norton and Ariely, *Perspectives on Psychological Science*, 6, 9–12) asked respondents to indicate what percent of the nation's total wealth is—and should be—controlled by richer and poorer quintiles of the U.S. population. We show that such measures lead to powerful anchoring effects that account for the otherwise remarkable findings that respondents reported perceiving, and desiring, extremely low inequality in wealth. We show that the same anchoring effects occur in other domains, namely web page popularity and school teacher salaries. We introduce logically equivalent questions about average levels of inequality that lead to more accurate responses. Finally, when we made respondents aware of the logical connection between the two measures, the majority said that typical responses to the average measures, indicating higher levels of inequality, better reflected their actual perceptions and preferences than did typical responses to percent measures.

Keywords: inequality, response bias, anchoring-and-adjustment, replication study.

1 Introduction

National differences in wealth and income inequality are large and important (Wilkinson & Pickett, 2010). But how much inequality do citizens tend to perceive? And how much inequality do they think there ought to be? The latter issue, preferences for how wealth, income and other valuable resources are distributed, has a long history in the social sciences and has been addressed by researchers in sociology (e.g., Cook & Hegtvedt, 1983), political science (e.g., Frohlich, Oppenheimer & Eavey, 1987), economics (e.g., Fehr & Schmidt, 1999), psychology (e.g., Eriksson & Simpson, 2011), and even biology (e.g., Brosnan & de Waal, 2003). The former issue, whether distributions are accurately perceived, has received much less attention but is arguably just as important. After all, many resource distributions, such as wealth and inequality in nations with millions of citizens are typically not directly observable. And preferences for redistributions are likely to be more influenced by how much inequality is perceived than how much actually exists. Yet, remarkably, we know relatively little about whether citizens of countries, such as the U.S., with highly unequal distributions of wealth and income actually perceive that inequality.

A recent paper by Norton and Ariely (2011) sought to address both how much wealth inequality Americans per-

ceive and how much inequality they think ought to exist. They found that Americans believe that the wealth distribution is dramatically lower than is actually the case, and prefer an even more equal distribution than the one they erroneously believe exists. These findings are surprising for a number of reasons, not least because they seem at odds with conventional wisdom. It is no secret that America's poor are very poor and its rich are very rich; further, political agendas for egalitarian distributions of wealth are relatively unpopular.

Here we ask whether the surprising findings reported by Norton and Ariely, rather than showing that Americans have remarkably poor understanding of the extent of U.S. wealth inequality and remarkably egalitarian preference, may stem from the particular measures of perceived and preferred inequality used by the researchers. Specifically, we investigate whether the surprising findings might depend on the original study's use of the following "Percent measure": *What percent of the United States' total wealth is [should be] controlled by the richest 20% of Americans?* (The same question was asked about all five quintiles down to the poorest 20%.) Note that the total wealth of quintiles is not directly observable but must be calculated as an aggregate of individual households. The computational demands of such a calculation may make many respondents unable to use their knowledge about rich and poor households to infer the total wealth of quintiles. To address this we adapted the measure to relieve respondents of the need to aggregate within quintiles. Our alternative "Average measure" instead simply asked respondents to indicate the average wealth of individual households: *What is [should be] the*

Work supported by the Swedish Research Council [2009-2390] and the National Science Foundation [SES-0551895, SES-0647169].

*School of Education, Culture and Communication, Mälardalen University, Box 883, SE-72123 Västerås, Sweden. Email: kimmo.eriksson@mdh.se.

[†]Centre for the Study of Cultural Evolution, Stockholm University.

[‡]University of South Carolina.

average household wealth, in dollars, among the 20% richest households in the United States?

Logically, the Percent and Average measures of inequality are intimately connected. For any pair of quintiles, the ratio between their percentages must be the same as the ratio between their averages. However, suppose respondents tend to be unable to make this logical connection. It may be that respondents tend to have reasonably accurate perceptions of wealth inequality, which they will express in responses to the Average measure, and yet give highly biased responses to the Percent measure. Indeed, bias in the latter case may arise from an anchoring-and-adjustment heuristic (e.g., Epley & Gilovich, 2006). Faced with the Percent measure, respondents may anchor on an equal distribution (20% per quintile) and make only a modest adjustment. Using this heuristic, people may respond with a distribution that they would not stand by if they were (conceptually and computationally able) to reflect more carefully about it.

Of course, inequality exists, at varying levels, in many domains. If our above arguments are correct, the Percent measure is predicted to yield low estimates of perceived inequality for any domain. The Average measure, on the other hand, is predicted to yield estimates that more accurately reflect the differences in levels of inequality between domains. In addition to the wealth domain we therefore studied *school teacher salaries* and *number of visits to web pages*, domains where inequality is known to be very low and very high, respectively.

2 Studies 1–4

We conducted four online surveys recruiting participants among users of the Amazon Mechanical Turk (Buhrmester et al., 2011; Paolacci et al., 2010). For each respondent, we used the Percent and Average measures to assess one of the following: ideal distribution of household wealth¹, or perceived distribution of household wealth, teacher salaries, or web page visits. The number of respondents who completed each survey are given in Table 1. This table also presents, for the bottom (poorest) and top (richest) quintiles, the mean value and standard error of responses to the Percent measure as well as the median value of responses to the Average measure (as extreme outliers biased the mean values).²

¹Following Norton and Ariely, wealth was defined to participants as follows: "Wealth, also known as net worth, is defined as the total value of everything someone owns minus any debt that he or she owes. So a person's net worth includes his or her bank account savings plus the value of other things such as property, stocks, bonds, art, collections, etc., minus the value of things like loans and mortgages."

²Following the procedure of Norton and Ariely (personal communication), for those few respondents who did not enter percentages that summed to 100, we rescaled the numbers to sum to 100; similarly, for those who entered numbers such that a quintile defined to be richer than

The next column gives the ratio between the just reported values for the top and bottom quintiles. An extreme ratio signifies high inequality. As we discuss below, the Percent measure and the Average measure tended to yield different ratios, i.e., different estimates of the inequality between the top quintile and the bottom quintile.

The bottom:top ratios for the Percent measure and the Average measure were also computed for each individual participant. The last column presents the proportion of respondents for whom this ratio in responses to the Average measure was more extreme (i.e., signified a higher estimate of inequality) than the ratio in responses to the Percent measure. A proportion higher than 50% is consistent with the Percent measure yielding responses biased towards lower inequality than elicited by the Average measure. As indicated in the table, all proportions were significantly different from 50% (related-samples sign tests, equivalent to binomial tests of the proportions being different to 50%) .

First consider the results for the ideal distribution of household wealth. The Percent measure yielded a ratio between the wealth of the poorest and richest quintiles of 1:4. This ratio is highly egalitarian and replicates the Norton-Ariely findings. But the Average measure yielded more than a twelve-fold increase of the ratio (1:50). This ideal distribution is certainly more egalitarian than the actual distribution of wealth but not nearly as dramatically egalitarian as the ideal given by the Percent measure.

The remaining surveys had respondents estimate perceived inequality. As predicted, the Percent measure yielded estimates of low inequality across all domains (ratios between 1:5 and 1:21), whereas the Average measure gave estimates of high inequality (on the order of 1:1,000) in those domains where inequality is, in fact, very high³, and a low estimate (1:2.4) in the low-inequality domain.⁴

3 Study 5

We conducted a fifth study that measured perceived and preferred wealth inequality using the Percent and Average measures. We then made respondents aware of the logical connection between the two measures and asked them to indicate whether typical responses to the Percent

another was actually poorer, numbers were reordered to correctly reflect the quintile order.

³The actual bottom:top ratio for wealth in the U.S. is on the order of 1:1000 (Norton & Ariely, 2011).

⁴Although peripheral to the purpose of this study, it is interesting that the Percent measure elicited *higher inequality* estimates than the Average measure for school teacher salaries. One interpretation is that the proposed anchoring-and-adjustment heuristic for quintile percent estimations do not actually anchor on the equal distribution but on some prototypical distribution of moderate inequality.

Table 1: Mean (\pm SE) response to the Percent measure and median response to the Average measure in estimations of the bottom and top quintiles in various domains.

Estimation task	N	Measure	Bottom quintile	Top quintile	Ratio	Lower inequality
Ideal household wealth	249	Percent	9.2(\pm 0.5)%	36.9(\pm 0.9)%	1:4.0	82%***
		Average	\$20,000	\$1,000,000	1:50	
Household wealth	570	Percent	3.1(\pm 0.2)%	65.0(\pm 0.8)%	1:21.2	75%***
		Average	\$1,000	\$1,500,000	1:1,500	
School teacher salaries	100	Percent	7.6(\pm 0.9)%	39.5(\pm 1.3)%	1:5.2	14%***
		Average	\$25,000	\$60,000	1:2.4	
Web page visits	100	Percent	5.5(\pm 0.4)%	47.4(\pm 1.6)%	1:8.6	72%***
		Average	100	110,000	1:1,100	

*** $p < .001$, related-samples sign test.

Note: The column Ratio gives the ratio between the group aggregate estimates for the bottom and top quintile (i.e., the previous two columns). The last column gives the proportion of respondents who estimated lower inequality between the top and bottom quintiles when they responded to the Percent measure than when they responded to the Average measure.

The data included some unlikely responses (e.g., negative teacher salaries in the bottom quintile); however, exclusion of such outliers does not change results in any substantive way.

or Average measures best reflected their perceptions and preferences for inequality.

3.1 Participants

Three hundred American participants (53% female; mean age 34 years, standard deviation 11 years) were recruited on the Amazon Mechanical Turk.

3.2 Materials

The survey had two parts. The first part measured respondents' perceptions of the level of wealth inequality, as well as how they thought wealth should be distributed (ideal) using both the Percent and Average elicitation methods. The second part of the survey made respondents aware of the logical connection between the elicitation methods and the inconsistency of their previous responses: "When we asked hundreds of people these questions, the typical answer was that the richest 20% currently control 65% of the total wealth and the poorest 20% currently control 3% of the total wealth, that is, a ratio of about 20:1. But when we asked about the current average wealth of the richest 20% and poorest 20%, typical answers were \$1,500,000 and \$1,000, respectively. These figures give a ratio of 1,500:1, which is obviously very different from the former ratio of 20:1."

Respondents were asked to consider whether their own responses had been inconsistent and, if so, why. A multiple choice format was used, with four options: (a) I did not really understand the questions at all; (b) I did not really understand there was a connection between categories' percentages of total wealth and the average wealth of individual households; (c) My answers were actually consistent; (d) Other reason (please specify).

The questionnaire then continued: "Given that you are now aware of the connection between percentage of total wealth and average household wealth, which of these answers best reflects your actual beliefs about how wealth currently is distributed?" A binary choice between a top:bottom quintile ratio of 1,500:1 and a ratio of 20:1 was given. Finally, the corresponding inconsistency for ideal inequality was presented. As in the previous question about perceptions of inequality, respondents were given a binary choice between 50:1 and 4:1 as the top:bottom quintile ratio that best reflected their actual beliefs about how wealth should be distributed.

3.3 Results

First, responses to the first part of the study closely replicated the results presented in Table 1. The Percent measure yielded 61.4(\pm 1.2)% vs. 3.1(\pm 0.3)% as mean es-

Table 2: Mean (\pm SE) response to the Percent measure and median response to the Average measure in estimations of the wealth of the bottom and top quintiles, broken down on self-reported level of comprehension.

Comprehension level	N	Question type	Bottom quintile	Top quintile	Ratio
Did not understand the questions	17	Percent	5.2(\pm 0.8)%	51.1(\pm 4.7)%	1:9.9
		Average	\$9,000	\$3,000,000	1:333
Did not understand the connection between the measures	174	Percent	3.0(\pm 0.4)%	61.1(\pm 1.5)%	1:20.5
		Average	\$5,000	\$1,000,000	1:200
Claimed to be consistent across measures	55	Percent	4.0(\pm 0.6)%	59.9(\pm 2.9)%	1:14.8
		Average	\$5,000	\$2,000,000	1:400
Others	44	Percent	1.9(\pm 0.3)%	68.7.4(\pm 2.9)%	1:36.0
		Average	\$5,000	\$1,750,000	1:350

*** $p < .001$, related-samples sign test.

Note: The column Ratio gives the ratio between the group aggregate estimates for the bottom and top quintile (i.e., the previous two columns).

imates of current wealth of the top and bottom quintiles, and 37.5(\pm 0.9)% vs. 9.4(\pm 0.3)% for ideal wealth. For the Average measure, the corresponding median responses were \$1,500,000 vs. \$6,000 for current wealth and \$1,000,000 vs. \$25,000 for ideal wealth.

Second, consistent with our arguments, a substantial majority of respondents, 58%, stated they had not understood there was a connection between percentage of total wealth owned by a given quintile and average wealth of individual households. Another 5% admitted to not having understood the questions at all. A minority, 18%, claimed to have given consistent answers to the two different measures—but, as shown in Table 2, they had actually been just as inconsistent as the others. Perhaps these participants incorrectly recalled their answers as having been consistent, but we think it is likely that their conflicting responses rather reflect poor comprehension of the entire set of questions.⁵

Finally, a majority of respondents chose the high inequality option—a top:bottom ratio of 50:1 rather than 4:1—as best reflecting their actual beliefs for ideal inequality (56.8%, $p = .023$, binomial test). For perceived inequality the tendency to choose the high inequality—a top:bottom ratio of 1,500:1 rather than 20:1—as best reflecting their actual beliefs was even stronger (70.2%, $p < .001$, binomial test). Thus, when made aware of the logical connection between the Percent measure and the Average measure, many people do not stand by the low inequality responses elicited by the Percent measure.

⁵The majority of respondents to Study 1 failed at least one of several comprehension checks of the Percent measure that were included in that study. Details available from the authors.

4 Discussion

Norton and Ariely (2011) set out to answer a remarkably important set of questions: How much inequality do ordinary Americans believe exists in the U.S.? And how much inequality do they desire? These questions have a range of important policy implications. However, the initial answers to these questions need to be reconsidered. Our findings indicate that the remarkably low estimates of wealth inequality given by Norton and Ariely’s respondents depended on the particular measure (quintile percentages) that was used. When asked to state quintile percentages for either household wealth, teacher salaries, or web page clicks, our respondents seemed to use an anchoring-and-adjustment heuristic leading to very similar responses across very different domains.

When respondents were relieved of aggregating their intuitions about inequality into quintile percentages, another picture emerged. According to this new picture, Americans do not tend to have extremely biased perceptions of current levels of inequality. Nor do they entertain an ideal of near-perfect egalitarianism. Rather they seem to prefer a world where the poor are not as poor as they are today. Further investigation into this more tractable ideal might provide a basis for workable policy prescriptions.

References

Brosnan, S. F., & De Waal, F. B. (2003). Monkeys reject unequal pay. *Nature*, 425, 297–299.

- Buhrmester, M. D., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6, 3–5.
- Cook, K. S., & Hegtvedt, K. A. (1983). Distributive justice, equity and equality. *Annual Review of Sociology*, 9, 217–241.
- Epley, N., & Gilovich, T. (2006). The anchoring and adjustment heuristic: Why adjustments are insufficient. *Psychological Science*, 17, 311–318.
- Eriksson, K. & Simpson, B. (2011). Perceptions of unfairness in allocations between multiple recipients. *Cognitive Psychology*, 62, 225–244.
- Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114, 817–868.
- Frohlich, N., Oppenheimer, J. A., & Eavey, C. L. (1987). Choices of principles of distributive justice in experimental groups. *American Journal of Political Science*, 31, 606–636.
- Norton, M., & Ariely, D. (2011). Consensus on building a better America—one wealth quintile at a time. *Perspectives on Psychological Science*, 6, 9–12.
- Paolacci, G., Chandler, J., & Ipeirotis, P. G. (2010). Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, 5, 411–419.
- Wilkinson, R., & Pickett, K. (2010). *The spirit level: Why equality is better for everyone*. London, England: Penguin.