**Citizens are not as misinformed as they first appear: Errors in estimates of
population size reflect common psychological processes**

D. Landy^a, B. Guay^b, T. Marghetis^a

Affiliations:

^a Department of Psychological and Brain Sciences, Indiana University

1101 E. 10th St., Bloomington, IN 47405

^b Department of Political Science, Duke University

140 Science Dr., 208 Gross Hall, Durham, NC 27708

Word counts:

Correspondence to:

Prof. David Landy

Psychological and Brain Sciences, Indiana University

1101 E. 10th St., Bloomington, IN 47405

e: dlandy@indiana.edu

t: 812-336-4449

Abstract

Successful democracies rely on citizens who act on the basis of the information available to them. If that information is incorrect, then democracy can suffer, since misinformed voters often make misinformed decisions. Unfortunately, when it comes to knowledge of demographic facts, misinformation appears to be the norm. Americans' tendency to massively overestimate the proportion of Americans who are immigrants, Muslim, LGBTQ, and Latino, but underestimate those who are white or Christian has been well documented and publicized over the past decade. Previous explanations have invoked issue-specific mechanisms such as xenophobia or media bias to explain these estimation errors. We reconsider this pattern of errors in the light of more than 30 years of research on the psychological processes involved in proportion estimation. In two publically available datasets featuring demographic estimates from 14 countries, we find that proportion estimates in macro-scale demographic situations correspond closely to what is found when people estimate much smaller objects in lab conditions, and when they make economic decisions. We conclude that demographic estimation biases are part of a very general pattern of human behavior—independent of the particular topic or demographic under consideration—and that this pattern explains most of the obvious error in estimations of politically salient population sizes. Topic-specific explanations such as media bias and xenophobia must be reinterpreted in light of a richer understanding of how humans estimate quantities generally.

Keywords: political misinformation, demographic estimation, probability weighting, social perception

1 Introduction

Justifications for democracy generally presume that, on average, citizens make decisions on the basis of realistic perceptions of the world (Fernbach et al, 2013; Gilens, 2001; Kuklinski et al, 2000; Somin, 2013, 2014). Unfortunately, across many different nations, communities appear to be massively distorted in their self-perception, often wildly misestimating the proportion of various demographic groups within their nation (Kuklinski et al, 2000; Citrin and Sides, 2008; Sigelman and Niemi, 2001; Lawrence and Sides, 2014), and similarly, though less powerfully, in local communities (e.g., Wong, 2007). For instance, a recent Gallup poll found that US residents estimate that over 20% of the population self-identifies as LGBT (Newport, 2015). The true proportion seems closer to 3% (Gates & Newport, 2012). Other populations, such as the proportion of US citizens who are white or Christian, are massively underestimated (e.g., Wong, 2007; Ipsos Social Research Institute, 2014). Similar biases are reported to affect estimates of the proportion of African Americans who are on welfare (Kuklinski et al, 2001), and the proportion of the total US population made up by Hispanics, Asian Americans (Wong, 2007), immigrants (Sides & Citrin, 2007), and African Americans (Sigelman & Niemi, 2001).

Since misinformed voters often make misinformed decisions, this phenomenon potentially has serious consequences for the optimal functioning of democratic systems (Somin, 2013, 2014). It is thus critical to understand the extent and root causes of apparent widespread misinformation. Recently, Ipsos Social Research Institute has taken a step towards understanding the extent of misinformation about the size of demographic groups by documenting patterns of misestimation of different sub-populations across countries (Ipsos Social Research Institute, 2014). These surveys provide a view unparalleled in breadth into a widespread social phenomenon: in countries around the world, people massively overestimate

the prevalence of minorities and small sub-populations, while dramatically underestimating the size of dominant groups.

Explanations of this phenomenon have invoked various topic-specific explanations. Some explanations posit ‘phobic innumeracy’ (Whaley & Link, 1998; Sides & Citrin, 2007; c.f. Martinez, Walden and Craig, 2008), in which people overestimate the prevalence of the things that they fear (e.g., immigrants). Other accounts speculate that estimates may be biased by media attention (Herda, 2013) or social contact (Herda, 2013; Wong, 2007), with individuals overestimating the size of groups that they are more exposed to.

The media has given extensive—often alarmist—coverage to these reports of apparent ignorance. For instance, in an article on demographic ignorance, a headline in *The Guardian* screamed, “Today’s key fact: you are probably wrong about almost everything” (October 29, 2014). Similarly, a piece for the news-magazine *Slate* had the pessimistic headline, “Americans Drastically Overestimate How Many Unauthorized Immigrants Are In The Country, And They Don't Want To Know The Truth” (January 9, 2012). More recently, the *Washington Post* chastised Americans for grossly overestimating the proportion of people in the US who are immigrants or Muslim: “But while many Americans consider immigration one of the biggest issues for the future president, [...] it’s remarkable just how much Americans overestimate immigration in their country. [...] American estimates for the size of the Muslim population in this country, also a focus of political discussion, are even more extreme” (September 1, 2016). Even the *Wall Street Journal* noted that, while “Americans have strong opinions about policy issues shaping the presidential campaign, from immigration to Social Security...their grasp of numbers that underlie those issues can be tenuous” (January 7, 2012). The mainstream consensus, therefore, is that people are often massively wrong when it comes to the demographic

facts that are relevant to critical political issues. And the origins of these errors? Journalists agree with academics: Issue-specific biases. Why do people overestimate the proportion of LGBT Americans? Bloomberg suggests it's because "gay, lesbian, bisexual, and even transgender characters have become prominent in recent years on TV shows such as 'Modern Family,' 'Scandal,' 'Degrassi,' and 'Glee,' as well as in movies" (May 22, 2015). Why do people overestimate the proportion of crimes committed by people of color? An op-ed in the *New York Times* argues that the cause is an underlying "racial bias" (September 7, 2014). It has become part of the public discussion, therefore, that personal and issue-specific biases have caused massive political ignorance, especially when it comes to politically-relevant demographic proportions.

These topic-specific explanations, however, cannot account for the striking regularities in misestimation across specific issues, populations, and surveys. Topic-specific explanations predict strong individual differences across country, topic, and time period. For instance, subpopulations which are overrepresented in the media or are seen as a cause of fear should be heavily overestimated, with more estimation error as media bias or fear increase. But if we zoom out to consider a wide variety demographic estimates, instead of focusing in on particular hot-button topics, a systematic pattern is glaringly obvious: small values are overestimated and large values are underestimated, *regardless of topic*. Previewing our results, Figure 2 shows the strength of this pattern, manifested as a clear S-shaped curve that cuts across a wide range of countries, issues, and proportions. While there is some question-by-question variability, it appears that most estimation error is driven by this S-shaped curve. One consequence of this S-shaped pattern is that the degree of overestimation is related systematically to the true size of the quantity being estimated: the smaller the true proportion, the more it is overestimated. This

striking regularity suggests that one, if not *the*, major cause of widespread demographic misestimation is not topic-specific biases, but some domain-general psychological process related to quantitative estimation.

This pattern of results is quite familiar to researchers who study the cognitive processes involved in quantitative estimation, particularly when those estimations involve some degree of uncertainty (Hollands & Dyre, 2000; Spence, 1990). Indeed, over the last several decades, this over-under pattern has been reported in fields as diverse as children searching through sandboxes (Huttenlocher, Newcombe, & Sandberg, 1994), the estimation of the age of the earth (Resnick, Shipley, Newcombe, Massey & Wills, 2012; Resnick & Shipley, 2013), numerical magnitudes (Barth & Paladino, 2012), proportions of dots presented on a screen that are white or black (Varey, 1990), and personal economic decision-making (Tversky & Kahneman, 1992). This recurring pattern of results is explained by a ‘ubiquitous’ property of human cognition: proportions are psychologically processed not as raw proportions, but as *log-odds* (Zhang & Maloney, 2012). The central claim of this paper, therefore, is that the interpretation of polls and surveys must take into account this domain-general feature of human proportional reasoning; only afterwards is there is any need to invoke theories about media bias, fear, or other social or informational factors.

In this paper, we first provide a brief review of psychological models of proportional reasoning in order to motivate intuitions for why proportional reasoning might function the way it does. We then discuss in detail a specific model of proportional reasoning, the *linear in log odds model*, that is especially well suited to judgments of demographic proportions. Next, we use this model to reanalyze two large, publically available polls in which people across 14 countries estimated the size of demographic groups within their countries. Our results suggest that most

demographic misestimation is a simple consequence of the quirks of human proportional reasoning, rather than evidence of issue-specific fear or media-motivated biases.

1.1 How do humans reason about proportions?

When people make judgments, they do not translate directly from any incoming information to explicit responses or reports. Instead, in general, there is a nonlinear relation between information and response (Fechner, 1860; Stevens, 1957; Spence, 1990). This insight is the basis for modern psychophysics. Psychophysics breaks down the route from initial information to explicit response into a series of steps. First, there is the raw, incoming information, which itself may be biased (e.g., in the context of demographic estimates, there may be media bias). Second, this raw information is *perceived* by the individual, which involves systematically applying a ‘perceptual function’ to transform incoming information into an underlying perception of the world. Third, the individual makes an explicit judgment on the basis of this underlying perception, with their response determined by a systematic ‘response function’ that transforms relevant perceptions into explicit responses.

To make sense of estimates of demographic proportions, we must consider the combined influence of the psychophysical transformations—the perceptual and response functions—that stand between the raw information in the world and an individual’s explicit judgments. Prior work on non-demographic proportion estimates suggests that these domain-general psychological processes can introduce reliable, systematic deviations. More specifically, when people estimate proportions, they very often overestimate small proportions, and underestimate large ones (Hollands & Dyre, 2000; Spence, 1990). A very similar phenomenon occurs with other kinds of quantitative estimates, such as probability judgments (Tversky & Kahneman, 1992; Gonzalez & Wu, 1999; Prelec, 1998).

Why might proportion estimation rely on psychological processes that introduce such reliably ‘wrong’ judgments? One possible reason is that, when judgments occur under uncertainty, it may be entirely rational to systematically overestimate small values and underestimate large values, since the judgment would then will reflect both the raw information and prior expectations about proportions in general (e.g., Huttenlocher et al, 1991; Fox and Rottenstreich, 2003; Fennell & Baddeley, 2012). If your incoming information—which might be visual, or based on memory, or from any combination of other sources—suggests an especially extreme value, then you might reasonably suppose that your sample of information must have been biased. Let’s say that you encounter a new species of tree, and the first three such trees that you encounter are exceptionally tall. What should you conclude about the average height of this species of tree? One approach would be to conclude, on the basis of that raw information, that this trees of this species, on average, grow very tall, much taller than the typical tree. Alternatively, one might conclude that those three individual trees must have been extreme outliers, since your past experience with average tree heights of other species suggests that such extreme heights are highly unlikely. Given the conflict between the heights that are *observed* and the heights that are *expected* to be seen, a reasonable thing to do is to split the difference and guess that trees of this new species are probably, on the whole, somewhat taller than other trees, though perhaps not as tall as one might conclude from the three observed trees alone. Similarly, if you were to encounter three trees of another new species, all three of which are exceptionally small, you might reach the complementary conclusion: perhaps this species of tree is short and scrubby, but likely somewhere between your small sample and the height of the typical tree that you’ve encountered (see Figure 1). More generally, evidence of any extreme value should rationally be attributed partially to a biased sample, and partially to a genuinely extreme value—

and, thus, explicit judgments should end up somewhere between the value suggested by the current information and the value suggested by prior beliefs. On Bayesian accounts of reasoning, this strategy—combining new information with prior beliefs—is a basic principle of rationality (Box & Tiao, 1973).

How do these considerations apply to estimates of demographic proportions? Here we provide two examples to illustrate the intuition of the psychophysical models presented below. Say that you are asked about the proportion of Americans who are Cambodian Americans. Sampling from your experience, you might expect this proportion to be quite low—perhaps less than 1%, if you have had low exposure to Cambodian people in your daily life. On the other hand, you have lots of experience with demographic subgroups in general. Most such groups are much larger than the Cambodian-American population, and only a very few are smaller¹. As in the case of the extreme heights, part of the extremity of this sample is probably due to the sample being small and noisy. Thus, on the basis of your experience with the distribution of demographic proportions, you could reasonably infer that your sample of Cambodian people is unrepresentative—that there are many more Cambodian people living in the US than your personal experience would suggest, but you just don't happen to encounter or recognize many Cambodian people. Even if your memory sample of Cambodian-Americans was actually unbiased—and, indeed, the true value is around 0.07%—then a rational reasoning process would

¹ Indeed, since for any demographic group (e.g., Asian Americans) one could be asked about its inverse (e.g., non-Asian Americans), the mean of demographic proportions is, in principle, exactly 50%—though it might well be that in actual experience one more often considers, or the media more often discusses, prominent minorities.

push you to systematically overestimate the proportion. Analogous considerations apply to majority sub-groups, which should be rationally underestimated. This process of rational adjustment can be conceptualized as a special case of regression to the mean, in which explicit estimates of demographic proportions are shifted systematically toward the mean proportion across all demographic groups (i.e., 50%). By taking into account previous knowledge, this strategy reduces overall error—but at the cost of introducing systematic biases, with smaller values overestimated and larger values underestimated (Box & Tiao, 1973; Huttenlocher et al, 1991). We will refer to this pattern as *uncertainty-based rescaling*, or simply *rescaling*.

Uncertainty-based rescaling is thought to be ubiquitous in quantitative judgments, and evidence for rescaling comes from a wide range of situations. Systematic overestimation of small values and underestimation of large ones has been observed in judgments in diverse domains, including number estimation (Landy, Charlesworth, and Ottmar, 2016; Cohen & Blanc-Goldhammer, 2011; Barth & Paladino, 2011), reading of bar graphs and pie charts (Spence, 1990), the perceived proportion of letters in a list of random letters (Erlick, 1964), the proportion of dots in a collection that are white or black (Varey, Mellers & Birnbaum, 1990), remembering the location of dots in space (Huttenlocher, Hedges, and Duncan, 1991), and risky events (Gonzalez and Wu, 1999; Tversky and Kahneman, 1992). Figure 1 presents a small sample of prior results.

Figure 1: Top panel: illustration of the rational process of adjusting extreme but uncertain perceptions toward aggregate prior expectations (i.e., a ‘typical’ value). Other panels: Representative results from prior research in proportion judgments, reprinted from Hollands & Dyre, 2000. Compare to Figure 2, which presents data specifically from demographic estimations. These studies cover a range of different psychophysical tasks: the proportion of presented dots that are of a particular color (Varey, 1990), the proportion of a time interval delimited by clicks (Nakajima, 1987), and the proportion of a presented sequence of letters that were the letter ‘A’. See original sources for details.

1.2 Modeling the psychology of proportion estimation

A number of different models have been developed to account for rescaling and other features of proportional and probabilistic reasoning. These include models founded on rational principles or cognitive models like the one described in the previous section (Huttenlocher, Hedges, & Duncan, 1991; Spence, 1990; Hollands & Dyre, 2000; Cohen, & Blanc-Goldhammer, 2011; Asano, Basieva, Khrennikov, Ohya, & Tanaka, in press) and models that provide good fits to particular datasets without particular justification (Tversky & Kahneman, 1992; Prelec, 1998). While these models differ in the details of their predictions, they all capture the same pattern of systematic over- and underestimation that typically appears in proportional reasoning in general, and in demographic estimates in particular. In the current study, to illustrate the importance of considering psychophysical scaling when interpreting the results of surveys asking about the size of demographic groups, we adopted an established model which lies somewhere in the middle: the *linear in log odds*, or LILO, model (Gonzalez & Wu, 1999). While this model does an exceptionally good job of accounting for demographic estimates, as we shall see, our intention is not to advocate for the LILO model in particular. Rather, we are using it to illustrate the more general point that apparent errors in demographic estimates might be nothing more than one manifestation of a ubiquitous property of human proportion estimation.

The LILO model was originally introduced as a descriptive model. The LILO model's descriptive success, however, likely stems from the fact that it formalizes three insights about the psychology of proportional reasoning (cf. Fox & Rottenstreich, 2003): (1) people encode proportions as *odds ratios*, (2) these odds ratios are represented, like other unbounded positive variables, on a *log scale*, and (3) people have prior expectations and uncertainties about

proportions. All three of these assumptions are supported by decades of research on the psychophysics of quantities.

For instance, decades of research has demonstrated that, under most circumstances, people process proportions not as percentages or probabilities, but as *odds ratios* (Fox & Rottenstreich, 2004; Gonzalez & Wu, 1999; Tversky & Kahneman, 1992; Shepard, 1981; Spence, 1990; Stevens, 1957; Zhang & Maloney, 2012). The difference is subtle but important. In a percentage, one considers the outcome of interest as a proportion of all possible outcomes (e.g., Asian-Americans / all Americans). In an odds ratio, one compares the outcome of interest with the *rest* of the possible outcomes (e.g., Asian-Americans / Non-Asian-Americans). For example, you may estimate that your train runs late about 2 out of every 10 times. You may then represent the proportion of ‘running late’ cases and consider it against the ‘on time’ cases, yielding an odds-ratio of .2 to .8, or 0.25. Odds ratios are rarely presented in scientific contexts, but are prototypical in betting situations. Notice that while comparison of odds ratios forms the basis of a variety of statistical techniques including logistic regression and risk analysis, here we consider simply the evaluation of a single odds ratio: the relation between the proportion of times that a thing happens and that it does not.

The second psychological motivation for the LILO model is that the mental representation of odds-ratios is log-scaled. This follows from the expanded range of odds ratios: for a proportion between 0 to 1, the associated odds ratio will be between 0 and infinity. The mental representation of most unbounded positive quantities is roughly log-scaled (Fechner, 1860; Dehaene, 2003; Shepard, Kilpatrick, and Cunningham, 1975; Zhang & Maloney, 2012), and it is unlikely that odds-ratios will be the exception to this psychological rule. On this assumption, the mental representation of a proportion, p , will consist of the log odds ratio, r :

$$(Eq. 1) \quad r = \log\left(\frac{p}{1-p}\right)$$

Conversely, to give an explicit estimate in terms of proportions, a log odds ratio r must be converted back into a proportion:

$$(Eq. 1') \quad p = \frac{e^r}{1+e^r}$$

One implication of this log scaling is that the psychological distance between proportions is given by their difference in log, rather than linear, space. For example, the psychological distance between a 20% chance of my train running late and a 50% chance should be roughly the same as that between a 20% chance and a 6% chance (i.e., since $\log(.5/.5) = 0$, $\log(.2/.8) \approx -1.4$, and $\log(.06/.94) \approx -2.8$).

The final underpinning of the LILO model is that people have prior expectations and uncertainty. People should expect, a priori, that an unknown demographic proportion will be somewhere in the vicinity of 0.5 (probably with quite wide variance), regardless of the specific demographic group being considered. People should take these expectations into account when evaluating their perceptions of the size of demographic groups. These expectations, moreover, involve a certain degree of uncertainty. This is typically implemented by assuming that expectations are normally distributed—in log-odds space, of course.

Taken together, these three insights about the psychological processing of proportions imply that we should *never* interpret explicit estimates as a direct reflection of underlying perceptions. Rather, estimates of demographic proportions should reflect a linear interpolation between underlying perceptions and a generic, domain-general prior, in log-odds space:

$$(Eq. 2) \quad \psi'(r) = \gamma r + \delta$$

in which r is the source proportion (i.e., the underlying perception), γ is the degree of bias toward the prior, δ is the location of the prior expectation (i.e., the belief about the value of a ‘typical’ proportion), and the result, $\psi'(r)$, is the predicted explicit response—all expressed as log-odds rather than as raw proportions. Since political surveys always require participants to respond using probabilities, not log-odds, this Eq. 2 can be transformed using Eqs. 1 and 1’ to give:

$$(Eq. 3) \quad \psi(p) = \frac{\delta p^\gamma}{\delta p^\gamma + (1-p)^\gamma}$$

This is equivalent to Eq. 2 except it is given in terms of the raw source proportion, p , instead of log-odds.

This model allows us to predict how a ‘perfect’ social observer would respond when asked to estimate demographic proportions. Even if somebody had correct, unbiased information about the exact, true values for all demographic proportions, their explicit estimates would exhibit systematic ‘errors,’ due solely to the psychological processes involved in proportional reasoning. Thus, if we assume that underlying perceptions are perfectly correct, then the LILO model can predict how a ‘perfect’ observer would transform their unbiased perceptions into explicit proportion estimates—estimates which, even for a ‘perfect’ observer, will reflect the rational rescaling that is characteristic of quantitative estimates.

1.3 Summary

Proportion estimation is an extremely well-studied phenomenon in human reasoning, and shows a consistent, clear pattern of inward bias toward mean or expected values. This has been documented across a wide variety of domains and activities. This has been ignored, however, when interpreting the public’s misperception of demographic populations. Discussions in political science and the media have not considered that demographic proportion estimation is

just a specific variety of proportion estimation more generally—and decades of research on proportion estimation has established that systematic errors are not always evidence of issue-specific misinformation. Therefore, when interpreting systematic errors in demographic proportion estimation, the first step should be to compare estimates to those of a ‘perfect’ observer who is subject to the basic, domain-general psychological processes that govern proportional reasoning. One way to do this is to apply an established psychophysical model of proportion estimation. One such model is the LILO model, which we shall use here.

This approach still leaves a role for previous explanations in the literature, such as media bias, disproportionate exposure to certain groups, or fear. Our point is that these issue-specific explanations should only be invoked after accounting for deviations that will occur naturally and rationally. Issue-specific explanations such as media bias, social contact, or xenophobia will likely be necessary to explain residual deviation between individuals’ estimates and the predicted estimates of a ‘perfect’ observer. But issue-specific explanations should *not* be invoked to explain errors in individual’s raw estimates. This is because perfectly rational, informed people will vastly overestimate low demographic proportions and underestimate high ones. The critical question, then, is whether past reports of widespread, issue-specific biases in demographic estimates are anything other than the rational rescaling that occurs during quantitative judgments under uncertainty.

2 Current Study

2.1 Method

To investigate whether considerations of domain-general psychological processing can account for previously reported demographic ignorance and misinformation, we reanalyzed two

large, multinational surveys for which the datasets are publicly available. Both these surveys have been used to argue for widespread topic-specific biases. The first was the *Ipsos MORI Perils of Perception Survey*, conducted in 2014 by the public policy polling center Ipsos Social Research Institute (Ipsos Social Research Institute, 2014). They polled individuals ($n = 11,527$) from 14 countries across North America, Europe, and Asia. Each individual was asked about the demographic proportions of their own country. The second was a subset of the *European Social Survey* (“ESS”), in which participants ($n = 38,339$) from twenty European countries were asked to estimate the proportion of their own country’s population that is foreign-born (Norwegian Social Science Data Services, 2006; analyzed previously in Sides and Citrin, 2007). Both studies reported massive errors in demographic estimates, and these errors have been taken as evidence of issue-specific bias (e.g., immigration is overestimated due to xenophobia).

The LILO model was hand-coded and all statistical analyses were conducted in the R statistical environment (R Core Team, 2015) using Eq. 3. Models were fit using maximum-likelihood methods, assuming normally distributed errors.

2.2 Results

Figure 2: Results of the 2014 Ipsos MORI *Perils of Perception* survey (A-B) and the European Social Survey (C-D). Left panels (A and C) present country-level mean estimates of a variety of demographic proportions. Note that smaller values are systematically over-estimated while larger values are systematically underestimated. The dashed red line indicates ‘perfect’ accuracy, where explicit estimates are equal to the true proportion. The grey line indicates predictions of a psychologically-plausible model of proportion estimation, the LILO model, assuming that individuals have unbiased perceptions of the true value but systematically adjust their explicit estimates toward a more ‘typical’ value. Right panels (B and D) present the same results but in terms of log odds-ratios. When converted to log-odds, the curve in Figures A and C becomes a straight line: the LILO model predicts a linear relation between the true log-odds and estimated log-odds. In these panels, the slope indicates the amount of rational rescaling toward a typical value (see Eq. 2).

For both polls, estimation errors were strongly predicted by the true value of the items being sampled: As expected, large values were systematically underestimated and small values overestimated, regardless of the specific topic or country (Ipsos: true for 95/98 estimated items, $p < .0001$, binomial test; ESS: 19/20, $p < .0001$). We compared these estimates to what we should expect if individuals had veridical, unbiased estimates of the true proportions, but rescaled all their estimates toward a generic expected value, using the LILO model. These predictions were an excellent fit to actual estimates, and could account for the observed relation between true values and raw estimation errors (Fig. 2, grey lines). Indeed, using the LILO model to account for the systematic rescaling predicted by psychological models of proportional reasoning had a dramatic impact on apparent bias. For the Ipsos poll, root mean squared error in the raw responses (RMSE) dropped from 17% without accounting for source-to-response rescaling, to just 6% after rescaling was taken into account. For the ESS, RMSE dropped from 9% to 5%. In other words, explicit estimates were actually quite close to what a ‘perfect’ observer might say, assuming that even a perfect observer will engage in rational rescaling when making explicit estimates. Most of the apparent ‘error’ in observed estimates is what we should expect if individuals are engaging in rational rescaling. Thus, error that has been interpreted previously as issue-specific “ignorance” is actually predicted systematically by a simple, issue-agnostic psychophysical model of proportion estimation under uncertainty.

Another way to visualize this result is in terms of log-odds-ratios, instead of raw proportions (Fig. 2B, D). The LILO model predicts a simple linear relation between true log-odds and estimated log-odds (cf., Eq. 2), with the slope of this relationship indicating the amount of rational rescaling toward a typical value. When interpreted as log-odds-ratios, explicit

estimates of demographic facts are related systematically and linearly to the true values of those facts.

Figure 3: Residual errors from the Ipsos MORI data set, separated by country and question.

Questions with higher cross-national objective values are higher. Blue bars indicate raw error; red bars indicate error after accounting for rescaling. Raw errors are both larger and more structured than post-model errors (that is, the questions with typically large values are underestimated, and small values overestimated).

If individual issues *were* subject to issue-specific biases such as xenophobia or media bias, this would manifest itself as reliable and systematic deviations *after* accounting for uncertainty-based rescaling. However, residual errors suggested patterns of bias that differed starkly from those indicated by previous work (Figure 3). In fact, for many items, the direction of residual bias after rescaling was in *the opposite direction* from the pre-rescaling error (Ipsos: 43/98; ESS: 9/20). For instance, every single country in the Ipsos poll overestimated the proportion of immigrants—sometimes by a factor of eight. But after accounting for rescaling due to uncertainty, most countries' estimates were actually *less* than the predicted estimate for proportions of that magnitude (i.e., most green dots lie below the grey line in Fig. 2A). In other words, after controlling for domain-general psychological processing of proportions, it appears that people may actually rely on information that underestimates the presence of immigrants in their country, relative to other demographic questions.

3 Discussion

The central thesis of this paper is simple: Demographic proportion estimation is a kind of proportion estimation, and is thus subject to the same psychological processes that humans use to estimate other proportions. But this uncontroversial observation has immediate and major implications for the interpretation of demographic misestimation. Decades of psychological research on proportional reasoning has established that explicit estimates are not direct reflections of underlying perceptions, but combine perceptions with prior expectations. As a result, surveys and polls that ask participants to estimate demographic proportions cannot be interpreted as direct measures of participants' underlying [mis-]information, since a large portion of apparent 'error' on any particular question will likely reflect rational, question-agnostic rescaling toward a more moderate expected value. Indeed, as predicted, the current study found

that a simple, domain-general model of proportion estimation can predict most of the ‘error’ in demographic estimates. Most of the difference between true values and explicit estimates appear to reflect rational rescaling under uncertainty, not issue-specific bias or misinformation.

The goal of the current study was to explore how explicit demographic estimates should deviate from the true value even if underlying perceptions of the world are unbiased; we thus assumed that people’s perceptions were normally distributed around the *true* value. After accounting for rational rescaling, residual errors were often in the opposite direction from what one would infer from the raw data; explaining these errors will thus require theoretical interpretations that reverse those that are dominant in the field. But the model can easily be adapted to account for issue-specific biases — such as those due to media bias or xenophobia — by assuming that perceptions are distributed around some other, biased value. Future work should investigate whether incorporating issue-specific biases can help account for the small residual errors that are not explained by domain-general, rational rescaling. Take media bias. It cannot account for the general pattern of over- and underestimation that we have documented here. But we can now ask whether, after accounting for psychophysical rescaling, media bias accounts for any of the residual error in individuals’ estimates.

While past work on political ignorance has ignored the implications of psychological rescaling, many social scientists already rely on statistical techniques that implicitly formalize the same central insight: that extreme values should be treated with suspicion and moved toward the expected value. Ridge or lasso regressions, for instance, involve a process known as *shrinkage* in which estimates of observed values are “shrunk” toward the expected value (e.g., Hastie, Tibshirani, & Friedman, 2009; Huttenlocher, Hedges, & Vevea, 2000). From a statistical perspective, this is an informed and even rational correction for unlikely values, and it is built

into many modern statistical inference processes, including some approaches to Bayesian statistical inference (Gelman & Shalizi, 2012). Therefore, when social scientists use these techniques, they are implicitly deploying one of the central principles of the LILO model. But what's good for the goose is good for the gander. If it is rational for social scientists to overestimate extremely low values and underestimate extremely high values, then it should be equally rational for everyday citizens, when faced with uncertainty about demographic proportions, to 'shrink' their extreme estimates toward a more moderate expected value.

3.1 Limitations to the LILO approach

There are some limitations to our analysis. First, as mentioned, there are a variety of psychological models of proportional reasoning that can account for the general pattern over- and underestimation that we observe here, and these make slightly different predictions. We do not claim that the LILO model is the best among those psychophysical models. (We note, however, that the LILO model provides an especially good fit to the data.) Our intention in using the LILO model here is only to illustrate the more general point that demographic estimates must be interpreted in light of the psychology of proportional reasoning. When interpreting any particular survey, poll, or study, *some psychological model* that accounts for rational rescaling is a necessary first analytical step before moving on to more issue-specific, ad-hoc explanations.

Second, the LILO model is applied here to group-level data, but was developed as a model of individual behavior. This is problematic, because the LILO model is not closed under averaging—that is, the mean of two individual log-linear response patterns is not usually itself log-linear in the aggregate. However, for the datasets analyzed in the current study, only mean responses were publicly available. While, typically, the average of many individual log-linear response patterns is itself a very close approximation to a log-linear response pattern, future

research should confirm the current results by analyzing responses on an individual level.

Indeed, ongoing work has found that patterns of over- and underestimation at the population level are well accounted for by LILO models of individual response patterns (citation to in prep study?).

Finally, it is worth noting that the LILO model was developed to capture the estimation of probabilities, rather than proportions. In principle, these may differ. However, it is reasonable to suppose that people estimate probabilities by considering the likelihood of each outcome, which would make the tasks essentially identical, as has been previously proposed (Zhang & Maloney, 2012). Furthermore, many tasks that have been taken as proportion judgment tasks can also be interpreted as probability estimates (e.g., Varey, 1990; Erlick, 1964), and the patterns of behavior that are observed in probability judgments are very, very similar to those observed in proportion judgments. Indeed, the LILO model of probability estimation is a slight generalization of a classic model of proportion estimation, the ‘power proportion’ model (Spence, 1990; Hollands & Dyre, 2000).

3.2 Future Research

Two outstanding analytical questions remain unanswered: First, is *individual* behavior also consistent with the psychophysical models? And do individual differences that predict differences in proportion estimation also predict demographic rescaling? For instance, the rational analysis presented here suggests that uncertainty should predict rescaling: the more uncertain I am about a true value, the more I should rely on my prior beliefs about the distribution of values typical to that category (e.g., if I see just one tall person in a country, I probably cannot confidently conclude that many people in that country are tall). Indeed, in other domains, uncertainty has been found to predict the tendency to rescale judgments toward an

expected or typical value (Huttenlocher et al, 1991; Landy, Crawford, & Salthouse, in press; Fox & Rottenstreich, 2003, Martins, 2006). The second question is how individual's demographic estimates relate to their political behavior. That is, do people vote, make policy decisions, or behave otherwise on the basis of their explicit estimates — which combine underlying perceptions with expectations — or do they behave on the basis of their underlying perceptions alone? In at least some decision-making contexts, rescaled probability estimates have been found to drive behavior in explicit tasks (Tversky & Kahneman, 1992). Alternatively, explicit estimates may only shape behavior when the task or situation requires individuals to construct an explicit estimate, thus prompting a rescaling of their implicit perceptions; in cases where explicit estimates are not salient, underlying perceptions may be more potent influences on behavior.

These results have important implications for policy makers, the press, and political scientists as well. Up to now, overestimation has been taken as evidence of personal or issue-specific biases—most commonly media bias, phobic innumeracy, or social contact biases. Our results indicate that overestimation is a natural consequence of making estimates about values that are smaller than the typical or expected value for members of that category (e.g., the proportion of the U.S. population that is Latino, which is smaller on average than most demographic groups). In order to conclude that there is a truly issue-specific bias, individuals' demographic estimates need to exceed what we should predict on the basis of rational psychophysical rescaling alone. In light of this, significant prior work must be reconsidered.

Finally, these results point to the need for greater collaboration between cognitive psychologists and psychophysicists on the one hand, and political scientists and polling agencies on the other. While political scientists are making increasing use of psychological theories and methodologies, political scientists and policy-relevant polling groups cannot be expected to

divine results and patterns learned over decades of research by psychologists. Simultaneously, psychologists are unlikely to understand the full ramifications their models and results in adjacent fields, like political science, that use similar methods to ask important and practical questions. Our results point to the potential for fruitful collaborations across disciplines, collaborations which can sharpen our responses to theoretical and empirical questions.

4 Author Contributions

DL and BG discussed theoretical developments in the topic of political innumeracy. DL conceived of the specific study idea and implemented the initial analysis. DL, TM, and BG considered together the implications. DL, TM, and BG wrote the paper.

5 Acknowledgments

Thanks to S. Attari, L. Nichols, and E. Weitnauer for feedback on various drafts of the manuscript, and to Simon Dedeo, Atrayee Mukherjee, and David Haussecker for useful conversations and comments.

6 References

- Asano, M., Basieva, I., Khrennikov, A., Ohya, M., Tanaka, Y. (in press). A quantum-like model of selection behavior. *Journal of Mathematical Psychology*.
<http://dx.doi.org/10.1016/j.jmp.2016.07.006>.
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68, 255-278.

- Barth, H. C., & Paladino, A. M. (2011). The development of numerical estimation: Evidence against a representational shift. *Developmental Science*, 14(1), 125–135. doi:10.1111/j.1467-7687.
- Box, G. E., & Tiao, G. C. (1973). *Bayesian Inference in Statistical Analysis*. Addison Wesley, Reading, Massachusetts.
- Citrin, J., & Sides, J. (2008). Immigration and the imagined community in Europe and the United States. *Political Studies*, 56, 33-56.
- Cohen, D. J., & Blanc-Goldhammer, D. (2011). Numerical bias in bounded and unbounded number line tasks. *Psychonomic Bulletin & Review*, 18(2), 331–338. doi:10.3758/s13423-011-0059-z
- Cokely, E. T., Galesic, M., Schulz, E., Ghazal, S., & Garcia-Retamero, R. (2012). Measuring risk literacy: The Berlin Numeracy Test. *Judgment & Decision Making*, 7, 25-47.
- Crawford, L. E., Landy, D., & Salthouse, T. A. (in press). Spatial working memory capacity predicts bias in estimates of location. *Journal of Experimental Psychology: Learning, Memory, & Cognition*. (available at davidlandy.net/wp-content/uploads/2015/11/Crawford-Landy-Salthouse-In-Press-2015.pdf).
- Dehaene, S. (2003). The neural basis of the Weber–Fechner law: A logarithmic mental number line. *Trends in Cognitive Sciences*, 7, 145–147. doi:10.1016/S1364-6613(03)00055-X
- Delli Carpini, M., & Keeter, S. (1996). *What the American public knows and why it matters* (Yale University Press, New Haven, CT)
- Erlick, D. E. (1964). Absolute judgments of discrete quantities randomly distributed over time. *Journal of Experimental Psychology*, 67, 475—482.

- Fechner, G. T. (1860). *Elemente der Psychophysik*. Leipzig: Breitkopf & Härtel
- Fernbach, P. M., Rogers, T., Fox, C. R., & Sloman, S. A. (2013). Political extremism is supported by an illusion of understanding. *Psychological Science*, 0956797612464058.
- Fennell, J., & Baddeley, R. (2012). Uncertainty plus prior equals rational bias: An intuitive Bayesian probability weighting function. *Psychological Review*, 119, 878-887.
- Fox, C. R., Rottenstreich, Y., (2003). Partition priming in judgment under uncertainty. *Psychological Science*, 14(3), 195-200.
- Gates, Gary J. & Newport, F (2012). Special Report: 3.4% of U.S. Adults Identify as LGBT. *Gallup*. Retrieved from <http://www.gallup.com/poll/158066/special-report-adults-identify-lgbt.aspx>.
- Gelman, A., & Shalizi, C. R. (2013). Philosophy and the practice of Bayesian statistics. *British Journal of Mathematical and Statistical Psychology*, 66(1), 8-38.
- Gilens, D. (2001). Political ignorance and collective policy preferences. *American Political Science Review*, 95, 379-396.
- Gonzalez, R., & Wu, G. (1999). On the shape of the probability weighting function. *Cognitive Psychology*, 38(1), 129-166.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Datamining, inference, and prediction* (2nd ed.). Berlin: Springer.
- Hollingworth, H. L. (1910). The central tendency of judgment. *Journal of Philosophy, Psychology & Scientific Methods*, 7, 461-469.

- Huttenlocher, J., Hedges, & L. V., & Duncan, S (1991). Categories and particulars: Prototype effects in estimating spatial location. *Psychological Review*, 98, 352–376.
- Huttenlocher, J., Hedges, & L. V., & Vevea, J. L. (2000). Why do categories affect stimulus judgments? *Journal of Experimental Psychology: General*, 129(2), 220-241.
- Huttenlocher, J., Newcombe, N., Sandberg, E. H. (1994). The coding of spatial location in young children. *Cognitive Psychology*, 27, 115-147.
- Ipsos Social Research Institute. (2014). *Ipsos MORI Perils of Perception Survey*.
<https://www.ipsos-mori.com/Assets/Docs/Polls/ipsos-mori-perils-of-perception-2014-tables.pdf>
- Kuklinski, J. F., Quirk, P. J., Jerit, J., Schwieder, D., & Rich, R-F. (2000). Misinformation and the currency of democratic citizenship. *Journal of Politics*, 62, 790-816.
- Landy, D., Charlesworth, A., & Ottmar, E. (in press). Categories of Large Numbers in Line Estimation. *Cognitive Science*, early online publication. 10.1111/cogs.12342
- Lawrence, E. D., & Sides, J. (2014). The consequences of political innumeracy. *Research & Politics*, 1(2), 2053168014545414.
- Lodge, M., & Taber, C. S. (2013). *The Rationalizing Voter*. (New York: Cambridge University Press)
- Lupia, A. (1994). Shortcuts versus encyclopedias: Information and voting behavior in California insurance reform elections. *American Political Science Review* 88, 63-76.

- Lupia, A. (2015). *Uninformed: Why People Seem to Know So Little about Politics and What We Can Do about It* (Oxford University Press, New York)
- Lynch, S. M. (2007). *Introduction to Applied Bayesian Statistics and Estimation for Social Scientists*. New York: Springer
- Martins A. C. R. (2006). Probability biases as Bayesian inference.judgment. *Decision Making 1*, 108–117.
- Martinez, M. D., Wald, K. D., Craig, S. C., (2008) Homophobic innumeracy? Estimating the size of the gay and lesbian population. *Public Opinion Quarterly* 72(4), 753-767.f
- Newport, F (2015). Americans Greatly Overestimate Percent Gay, Lesbian in U.S. *Gallup*. Retrieved from <http://www.gallup.com/poll/183383/americans-greatly-overestimate-percent-gay-lesbian.aspx>.
- Norwegian Social Science Data Services. (2006). *European Social Survey Round 3, Version 3.5*. [dataset] <http://www.europeansocialsurvey.org/data/download.html?r=3>
- Petrova, D. G., Pligt, J., & Garcia-Retamero, R. (2014). Feeling the numbers: On the interplay between risk, affect, and numeracy. *Journal of Behavioral Decision Making*, 27, 191-199.
- Prelec, D. (1998). The probability weighting function. *Econometrica*, 60, 497–528.
- R Core Team (2015). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Resnick, I. & Shipley, T.F. (2013). Application of the Category Adjustment Model in Temporal, Spatial, and Abstract Magnitude at the Billions Scale. In Knauff, M., Pauen, N., Sebanz

& I. Wachsmuth (Eds.), Proceedings of the 35th Annual Meeting of the Cognitive Science Society.

- Resnick, I., Shipley, T., Newcombe, N., Massey, C., Wills, T. (2012). Examining the Representation and Understanding of Large Magnitudes Using the Hierarchical Alignment model of Analogical Reasoning. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the 34th Annual Conference of the Cognitive Science Society*.
- Satopää, V. A., Baron, J., Foster, D. P., Mellers, B. A., Tetlock, P. E., & Ungar, L. H (2014). Combining multiple probability predictions using a simple logit model. *International Journal of Forecasting*, 30, 344-356.
- See, K. E, Fox, C. R., & Rottenstreich, Y. S. (2006). Between ignorance and truth: Partition dependence and learning in judgment under uncertainty. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 32(6), 1385-1402.
- Shepard, R. N. (1981). Psychological relations and psychophysical scales: On the status of “direct” psychophysical measurement. *Journal of Mathematical Psychology*, 24, 21–57.
- Shepard, R. N., Kilpatrick, D. W., & Cunningham, J. P. (1975). The internal representation of numbers. *Cognitive Psychology*, 7, 82-138.
- Sides, J., & Citrin, J. (2007). European opinion about immigration: The role of identities, interests and information. *British Journal of Political Science* 37, 477-504.
- Sigelman, L., Niemi, R-G. (2001). Innumeracy about minority populations: African-Americans and Whites compared. *Public Opinion Quarterly*, 65, 86-94.
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2012). A 21 word solution.
<http://ssrn.com/abstract=2160588>

- Somin, I. (2013). *Democracy and political ignorance: Why smaller government is smarter* (Stanford University Press, Stanford, CA)
- Somin, I. (2014). Why political ignorance undermines the wisdom of the many. *Critical Review*, 26, 151-169.
- Spence, I. (1990). Visual psychophysics of simple graphical elements. *Journal Experimental Psychology: Human Perception & Performance*, 16, 683-692.
- Stevens, S. S. (1957). On the psychophysical law. *Psychological Review*, 64, 153-181.
- Tversky, A., Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk & Uncertainty* 5, 297-323.
- Varey, C. A., Mellers, B. A., & Bimbaum, M. H. (1990). Judgments of proportion. *Journal of Experimental Psychology: Human Perception and Performance* 16, 613-625.
- Whaley, A L., and Link, B. G. (1998). Racial Categorization and Stereotype-Based Judgments about Homeless People, *Journal of Applied Social Psychology* 28,189–205.
- Wong, C. J. (2007). “Little” and “big” pictures in our heads: Race, local context, and innumeracy about racial groups in the United States. *Public Opinion Quarterly*, 71, 392-412.
- Zhang, H., & Maloney L. T. (2012). Ubiquitous Log Odds: A Common Representation of Probability and Frequency Distortion in Perception, Action, and Cognition. *Frontiers in Neuroscience*, 6(1). 10.3389/fnins.2012.00001

Draft