Bias and ignorance in demographic perception

D. Landy[a], B. Guay[b], T. Marghetis[a]

[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

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

When it comes to knowledge of demographic facts, misinformation appears to be the norm. Americans massively overestimate the proportion of their fellow citizens who are immigrants, Muslim, LGBTQ, and Latino, but underestimate those who are white or Christian. Previous explanations of these estimation errors have invoked topic-specific mechanisms such as xenophobia or media bias. We reconsider this pattern of errors in the light of more than 30 years of research on the psychological processes involved in proportion estimation and decision-making under uncertainty. In two publically available datasets featuring demographic estimates from 14 countries, we find that proportion estimates of national demographics correspond closely to what is found in laboratory studies of quantitative estimates more generally. Biases in demographic estimation, therefore, are part of a very general pattern of human psychology — independent of the particular topic or demographic under consideration — which explains most of the error in estimates of the size of politically salient populations. By situating demographic estimates within a broader understanding of general quantity estimation, these results demand a re-evaluation of topic-specific misinformation about demographic facts, along with topic-specific explanations of demographic ignorance such as media bias and xenophobia.

Keywords: judgment and decision making; political misinformation, demographic estimation, Bayesian modeling, social perception

1   Introduction

Justifications for democracy generally presume that the average citizen should make decisions on the basis of realistic perceptions of the world (see e.g., Fernbach et al, 2013; Gilens, 2001; Kuklinski et al, 2000; Somin, 2013, 2014). When considering policies like immigration, welfare, or social security reform, citizens should be expected to have some understanding of how many Americans are undocumented, receive welfare, or are eligible to receive social security benefits, respectively. Such perceptions are typically measured by asking people to make quantitative estimates of the prevalence of particular demographic groups, such as proportions of adherents to particular ethnic groups, social classes, or religions. Unfortunately, people across many different nations appear to hold massively distorted perceptions of the demographic composition of their communities, often wildly misestimating the proportion of various demographic groups within their local communities (e.g., Wong, 2007) and their larger nation (Kuklinski et al, 2000; Citrin and Sides, 2008; Sigelman and Niemi, 2001; Lawrence and Sides, 2014). For instance, a recent Gallup poll found that US residents estimate that over 20% of the population self-identifies as LGBT (Newport, 2015), but the true proportion seems closer to 3% (Gates & Newport, 2012). A similar bias toward overestimation is found in estimates of politically salient subpopulations — such as the proportion of African Americans who are on welfare (Kuklinski et al, 2001) — and in estimates of the size of minority groups, including Hispanics, Asian Americans (Wong, 2007), immigrants (Sides & Citrin, 2007), and African Americans (Sigelman & Niemi, 2001). This pattern reverses when people estimate the size of dominant groups, including US citizens who are white or Christian, which are massively underestimated (e.g., Wong, 2007; Ipsos Social Research Institute, 2014).

Since misinformed voters often make misinformed decisions, this phenomenon has potentially serious consequences for the optimal functioning of democratic systems of government (Somin, 2013, 2014, but see Lupia & McCubbins, 1998). It is thus critical to understand the extent and root causes of apparent widespread misinformation. As a step towards understanding the extent of misinformation about the size of demographic groups, a number of major international surveys have documented patterns of misestimation of different sub-populations within and across countries (e.g., Ipsos Social Research Institute, 2014; Norwegian Social Science Data Services, 2006). These surveys provide a view unparalleled in breadth into a widespread social phenomenon: in countries around the world, people massively overestimate the size of minority groups while dramatically underestimating the size of majority groups.

Most explanations of this phenomenon have focused on properties of the individual topic or population being estimated. One such explanation is that people overestimate the prevalence of the things they fear (e.g., immigrants), resulting in a form of 'phobic innumeracy' (Allport, 1954; Gallagher, 2003; Herda, 2010; Nadeau & Niemi, 1995; Gallagher, 2003; Whaley & Link, 1998). While this approach is not without its critics (Martinez, Walden and Craig, 2008; Sides 2016; Sides & Citrin, 2007), it illustrates the general tendency of the literature on demographic estimation to invoke explanations centered on topic-specific biases. Other explanations speculate that individuals overestimate the size of the specific groups to which they are overexposed, so that beliefs about population size are inflated by media over-representation (Gallagher, 2003; Herda, 2013) or increased social contact (Alba, Rumbaut, & Marotz, 2005; Herda, 2013; Wong, 2007).

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, like many scholars, assume that misestimation is a telltale sign of topic-specific bias. 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 a targeted 'racial bias' that specifically distorts beliefs about people of color (September 7, 2014). It has become part of the public discussion, therefore, that personal and topic-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 in misestimation patterns across country, topic, and time period. For instance, those specific subpopulations that are seen as a cause of fear or are over-represented in the media should be heavily overestimated. As the demographic groups that are feared or over-represented in the media differ from across countries, so should biases in estimation. But if we zoom out to consider a wide variety of demographic estimates across countries, instead of focusing only on a limited sample of hot-button topics within the U.S., 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 (inverted) S-shaped curve that cuts across a wide range of countries, topics, and values. 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 some domain-general psychological process related to estimation of demographic proportions.

Indeed, researchers working on quantitative estimation are familiar with this pattern of errors—overestimation of small values and underestimation of large values—particularly when those estimations involve some degree of uncertainty (Hollands & Dyre, 2000; Spence, 1990; cf. Gescheider,1976; Stevens, 1975). Indeed, 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), proportions of letters in a string that are 'a' (Erlick, 1964), and in personal economic decision-making as analyzed in prospect theory (Tversky & Kahneman, 1992). One upshot of this research is that perceptions can, and often *do*, differ starkly from their explicit expressions (Stevens, 1946; Mosteller & Tukey, 1977; Gescheider, 1976), since the process of transforming perceptions into responses can introduce systematic distortions. Nevertheless, research in political science often uncritically takes errors in estimation to be direct indicators of biased or misinformed perceptions.

The central claim of this paper is that the interpretation of polls and surveys must account for domain-general features of human proportional reasoning; only afterwards is there is any need to invoke issue-specific theories about media bias, fear, or other social or informational factors. In other words, before invoking domain-specific biases — such as homophobia or xenophobia — to explain demographic misestimation, one should first consider how a perfectly informed individual, with access to unbiased information about various populations, would estimate the relative sizes of those populations. This individual would be subject to the standard psychophysical mechanisms that are known from other instances of proportional reasoning — mechanisms which have the potential to make the individual appear misinformed when completing a survey. This approach is in line with prior work in political science that has sometimes suggested that general numeracy and cognitive function may be important for demographic estimation (Lawrence & Sides, 2014; Lunkmark & Kokkonen, 2014; Herda, 2015; Alba, Rumbaut, & Marotz, 2005; Kunovich, 2012). We expand on this by providing precise

predictive models of the connection between demographic perception, general numeracy, and explicit demographic estimates, grounded in the prior literature on psychophysical transformations.

In what follows, 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. In particular, we discuss two properties of human cognition that suffice to explain the S-shaped pattern of responses that is found in proportion estimation: (1) proportions are psychologically processed not as raw proportions, but as *log-odds* (Zhang & Maloney, 2012); and (2) human reasoning tends to combine new information with prior expectations (i.e., cognition is often "Bayesian," Huttenlocher et al, 1991; Lee & Danilieko, 2014). We then introduce a specific model of proportional reasoning that formalizes these two principles. 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 judgments, rather than evidence of topic-specific fear or media-motivated biases. While topic-specific biases may exist, most of the evidence that has been invoked as support for topic-specific biases is better explained by the domain-general psychological processes translate psychophysical stimuli into explicit estimates.

2   How do humans reason about proportions?

Human judgments are not a direct translation of incoming information and beliefs to explicit responses. Instead, in general, the psychological processes involved in representing and processing information introduce a nonlinear relation between information and response

(Fechner, 1860; Stevens, 1957; Spence, 1990). This insight is the basis of 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 about the world, which itself may be biased (e.g., in the context of demographic estimates, there may be media bias in the representation of the immigrant population). Second, this raw information is *perceived* by the individual, which involves processing this incoming information to create a perception of the world (e.g., a perception of the size of the immigrant population). Third, if an individual is asked to give an explicit response (e.g., a numerical estimate on a survey, voting, etc.), they must transform their perception from an internal scale into an explicit response. When an explicit estimate is incorrect, therefore, it could reflect any stage of the psychological process: bias in the raw information, bias introduced when creating a perception from this raw information, bias introduced when translating an perception into an explicit response, or any composition of these.

To make sense of the public's misestimation of demographic proportions, we must thus consider the combined influence of the psychophysical transformation from raw information to perceptions of the world, and then from these perceptions to explicit judgments. These psychological processes are not specific to demographic estimation. They occur whenever humans make judgments on the basis of information—any information, whether visual, based on memory, or from any combination of sources. Critically, prior work on non-demographic proportion estimates has demonstrated that these domain-general psychological processes introduce reliable, systematic deviations. When people generate explicit estimates of proportions, they 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 as studied in prospect theory (Tversky & Kahneman, 1992;

Gonzalez & Wu, 1999; Prelec, 1998).

Why might proportion estimation rely on psychological processes that introduce such

reliably 'incorrect' judgments? One possible reason is that, under many circumstances, it is

reasonable for an individual to rely not only on raw information, but also on prior expectations

about proportions in general. If an individual's information suggests an especially extreme value,

then they might suppose, reasonably, that the sample of information they have been exposed to

must have been biased. Indeed, when judgments occur under uncertainty, it can be rational to

combine raw information with prior expectations — even though this process can produce

explicit estimates that systematically overestimate small values and underestimate large values

(e.g., Huttenlocher et al, 1991; Fox and Rottenstreich, 2003; Fennell & Baddeley, 2012;

Petzschner, 2012). One good example of this kind of approach is the 'Decision by Sampling'

framework (Stewart, Chater, & Brown, 2006), according to which decisions are made by

combining information sampled from the decision context with beliefs about the background

distribution of relevant values. Systematic over- and underestimation are thus natural predictions

of both optimal and descriptive accounts of judgment and decision-making. On all these

accounts, information from the current context is combined with past beliefs, experiences, and

perceptions in ways that can manifest as systematic biases.

An example might help to illustrate the rationality of this general strategy. 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 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 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: trees of this species are small but likely somewhere between your small sample and the typical height of trees that you have 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 the prior distribution. On Bayesian accounts of reasoning, this strategy—combining newly sampled information with prior expectations—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 very few are smaller[1]. As in the case of the extreme tree 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 have not encountered or recognized them. Even if your memory sample of Cambodian-Americans was actually unbiased—and, indeed, the true value is quite small, around 0.07%—then a *rational* reasoning process would 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

---

[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.

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.

A second critical feature of proportional reasoning is that, under most circumstances,

people process proportions not as percentages or probabilities, but as *odds* (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). To calculate the odds, 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

of .2 to .8, or 0.25. Odds are rarely presented in scientific contexts, but are prototypical in betting

situations. Notice that while comparison of odds 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: the relation between the proportion of times that a thing happens and that it does

not. Moreover, since the mental representation of most unbounded positive quantities is roughly

log-scaled (Fechner, 1860; Dehaene, 2003; Shepard, Kilpatric, and Cunningham, 1975; Zhang &

Maloney, 2012), the mental representation of odds is likely also log-scaled. Thus, in accounting

for proportional reasoning, we must take into account the principle that the psychological

processes will operate over the relevant *log-odds,* not the raw proportions.
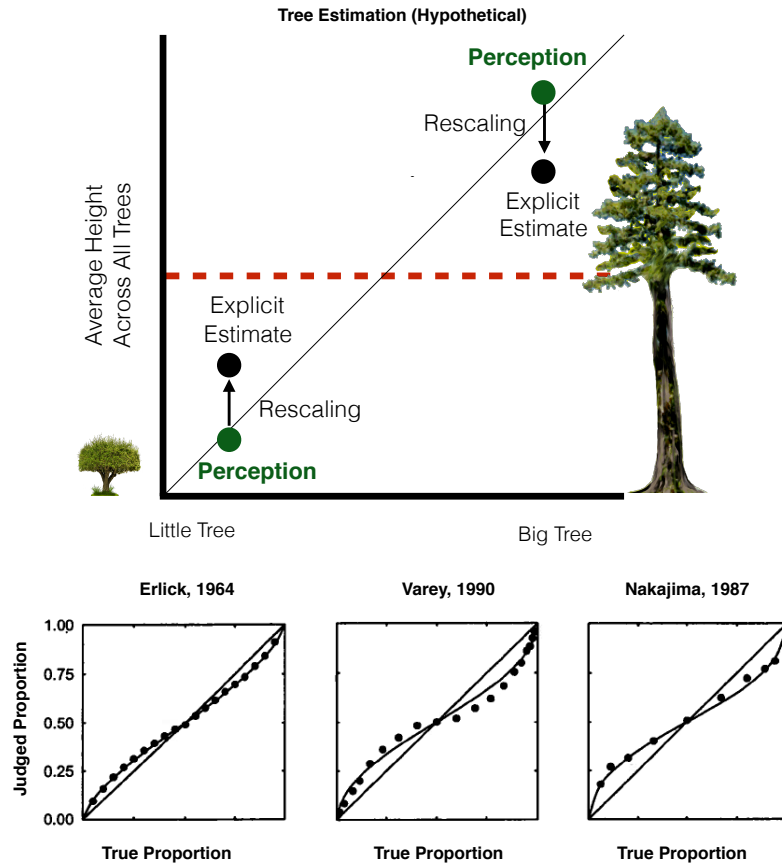

Figure 1: Top panel: illustration of the rational process of adjusting extreme but uncertain perceptions toward aggregate prior expectations (i.e., a 'typical' value). Other panels: Representative results from prior research in proportion judgments, reprinted from Hollands & Dyre, 2000. Compare to Figure 2, which presents data specifically from demographic estimations. These studies cover a range of different psychophysical tasks: (left) proportion of a presented sequence of letters that were the letter 'A' (Erlick, 1964); (middle) proportion of presented dots that are of a particular color (Varey, 1990); (right) proportion of a time interval delimited by clicks (Nakajima, 1987).

2.1   Modeling the psychology of proportion estimation

A number of different models have been developed to account for rescaling and other features of proportional reasoning (Huttenlocher, Hedges, & Duncan, 1991; Spence, 1990; Hollands & Dyre, 2000; Cohen, & Blanc-Goldhammer, 2011; Asano, Basieva, Khrennikov, Ohya, & Tanaka, in press; Lee & Danileiko, 2014; Tversky & Kahneman, 1992; Prelec, 1998). While these 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.

We sought a model that formalized both insights about the psychology of proportional reasoning: (1) people encode proportions as *odds*, and these odds are represented, like other unbounded positive variables, on a *log scale,* and (2) people have prior expectations and uncertainties about proportions, and they incorporate those priors into their explicit estimates (i.e., uncertainty-based rescaling). The first insight implies that the mental representation of a proportion, p, should consist of the log odds, r_p:

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

Conversely, to give an explicit estimate in terms of proportions, the (possibly transformed) log odds ($r_{p'}$) must be converted back into a proportion:

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

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 second insight — uncertainty-based rescaling — implies that estimates of demographic proportions should reflect an interpolation between perceptions and a generic, domain-general prior:

(Eq. 2) $$\psi'(r_p) = \gamma r_p + (1 - \gamma)\delta$$

in which r_p is the source proportion as perceived in log odds space, γ is the relative weighting of the perception and the prior, δ is the location of the prior expectation (i.e., the distribution of values of a 'typical' proportion), and the result, $\psi'(r_p)$, is the predicted explicit response—all expressed as log-odds rather than as raw proportions. Since political surveys require participants to respond using proportions, not log-odds, Eq. 2 can be transformed using Eqs. 1 and 1' to give:

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

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

This model allows us to predict how a 'perfect' 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 should exhibit systematic 'errors' due solely to the psychological processes involved in proportional reasoning. Thus, if we assume that perceptions are perfectly correct, then we 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.

In fact, Eq. 3 has been used previously, sometimes with slight modifications, to account for judgments about a variety of more mundane, human-scale proportions, preferences, probabilities, and risks (Birnbaum & McIntosh, 1996; Goldstein & Einhorn, 1987; Gonzalez & Wu, 1999; Karmarkar, 1978; Lattimore, Baker, & Witte, 1992; Tversky & Fox, 1995; c.f. Tversky & Kahnemann, 1992). We should note, however, that our intention is not to advocate for this particular model — although, as we shall see, it does an excellent job of predicting errors in demographic estimates. Rather, our intention is to illustrate the more general point that apparent errors in demographic estimates might be nothing more than one manifestation of ubiquitous properties of human proportion estimation under uncertainty.

2.2 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. But discussions in political science and the media have not considered that, as a specific variety of proportion estimation more generally, estimates of demographic proportions will be subject to the same domain-general sources of systematic error, which are thus not always evidence of topic-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 that formalizes the basic principles that govern proportion judgments under uncertainty.

This approach still leaves a role for previous explanations in the literature, such as media bias, disproportionate exposure to certain groups, or fear. But if domain-general psychological processing is responsible for most estimation error, then these topic-specific explanations should

only be invoked after accounting for deviations that will occur naturally and rationally. Topic-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 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, topic-specific biases in demographic estimates are anything other than the rational rescaling that occurs during quantitative judgments under uncertainty.

3   Current Study

3.1   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 sample was a non-probability based sample, which was weighted to match national demographics. 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 polls phrased questions in units of 100, e.g., "out of every 100 people living in [Country], how many immigrated to that country"? For details on sampling design, response rate, etc., please see the Appendix.

The model given by Eq. 3 was hand-coded and all statistical analyses were conducted in the R statistical environment (R Core Team, 2015). Models were fit using maximum-likelihood methods, assuming normally distributed errors. The full model and data set is available at https://osf.io/kt8wa/.

3.2   Results

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, as predicted by the 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 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 uncertainty-based 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 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 topic-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, instead of raw proportions (Fig. 2B, D). As illustrated by Eq. 3, the model predicts a simple linear relation between true log-odds and estimated log-odds with the slope of this relationship indicating the amount of rational rescaling toward a typical value. When interpreted as log-odds, explicit estimates of demographic facts are related systematically and linearly to the true values of those facts.

If individual issues *were* subject to topic-specific biases such as phobic threat 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 immigration-specific factors may be driving estimates of these populations *down*, not up, relative to other demographic groups of the same size. It is this residual error left after accounting for general psychophysical rescaling that could be explained by immigration-specific factors — factors such as underrepresentation in popular media and in many social networks.
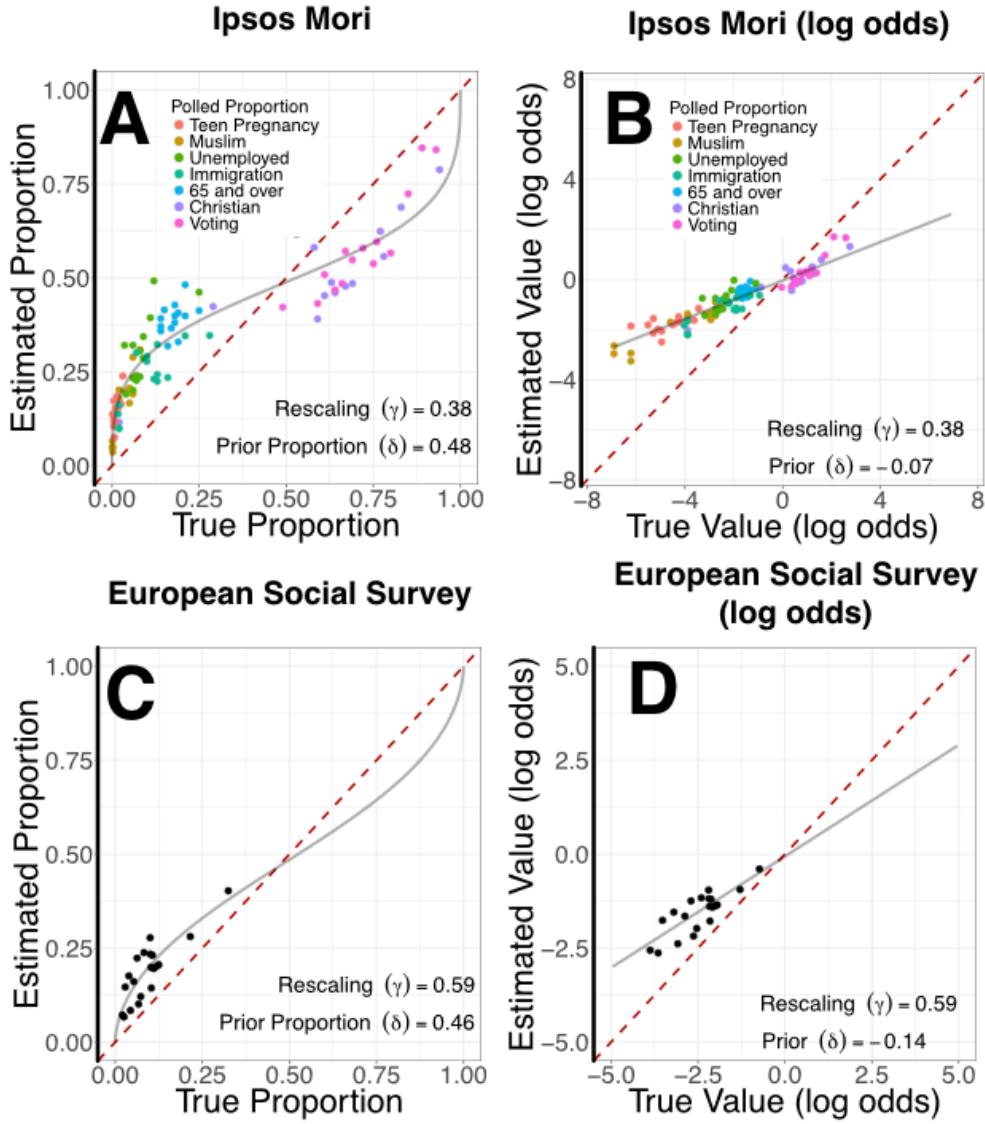
Figure 2: Results of the 2014 Ipsos MORI Perils of Perception survey (A-B) and the European Social Survey (C-D). Left panels (A and C) present country-level mean estimates of a variety of demographic proportions. Note that smaller values are systematically over-estimated while larger values are systematically 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 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. When converted to log-odds, the curve in Figures A and C becomes a straight line: the 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). Delta, the log-odds of the fitted prior, is given both as a proportion (left) and in log-odds space (right)..
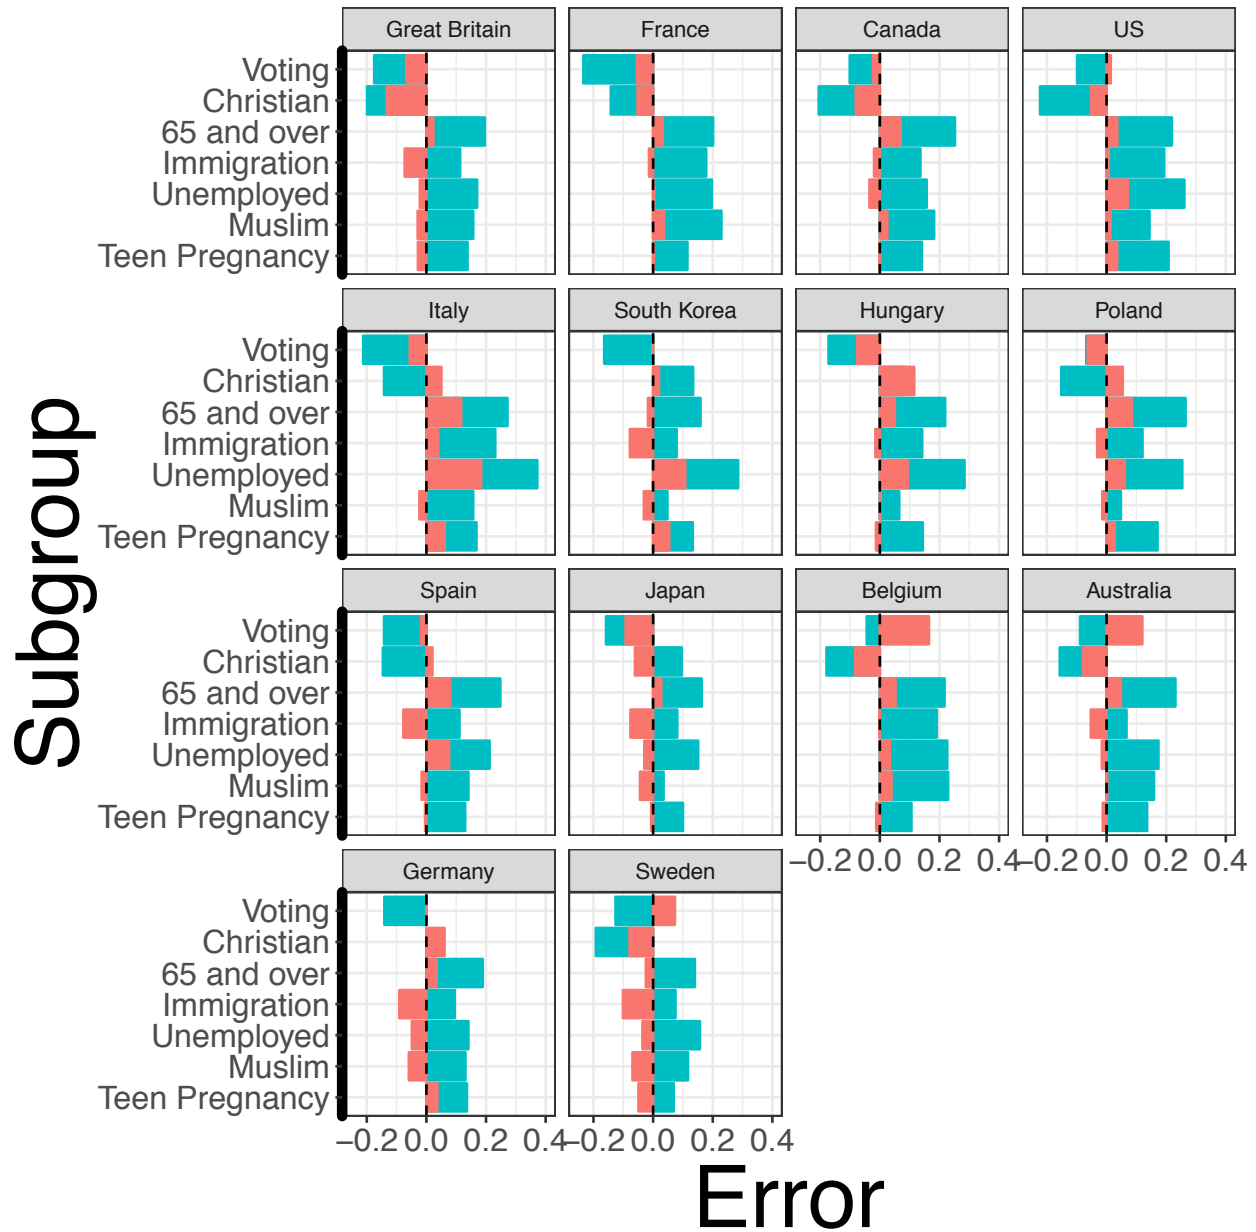
Figure 3: Residual errors from the Ipsos MORI data set, separated by country and question. The y-axis is ordered by the mean cross-national average objective proportion of the subgroup. Blue bars indicate raw error; red bars indicate error after accounting for rescaling. Raw errors are both larger and more structured than post-model errors (that is, the subgroups with typically large values are underestimated, and small values overestimated).

4   Discussion

Our central thesis is simple: Demographic proportion estimation is just one specific kind of proportion estimation, and thus subject to the same psychological processes that humans use to estimate other proportions. This banal 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 perceptions, but systematic transformations of those perceptions. As a result, surveys and polls that ask participants to estimate demographic proportions cannot be interpreted as direct measures of participants' [mis-]information, since a large portion of apparent 'error' on any particular question will likely reflect rescaling toward a more moderate expected value, regardless of the specific demographic population being estimated. 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. Of course, this work does not undermine the broader conclusion that citizens are, in general, distressingly misinformed about the world — for instance, about non-quantitative judgments (e.g., Nyhan & Reifler 2010, Berinsky 2015; Somin, 2013). When it comes to misinformation in quantitative estimates of demographic proportions, however, our results indicate that most of the difference between true values and explicit estimates reflects rational rescaling under uncertainty, not topic-specific bias or misinformation..

The two data sets examined here both show massive errors in demographic estimates, errors that have been taken as evidence of topic-specific bias or misinformation. For instance, Sides and Citrin (2007) state that "[i]f correct information about immigrant stock and flows reached the general public, our analysis suggests that the sense of 'threat' might wane, mitigating hostility towards immigrants." This interpretation makes sense only if people are, indeed,

misinformed about immigrant stocks and flows. After accounting for rational rescaling, however, residual errors in estimates of immigration were small and often in the opposite direction from what one would infer from the raw data. On our analysis, therefore, global overestimation of immigration follows naturally from general psychophysical biases; the puzzle that remains for topic-specific explanations is why immigration was not overestimated even *more*. Explaining these errors will thus require theoretical accounts that go in the opposite direction from those than are dominant in the field.

The goal of the current study was to explore how explicit demographic estimates would deviate from the true value if perceptions of the world were unbiased; we thus assumed that people's perceptions were normally distributed around the *true* value. However, in reality, individuals are unlikely to be perfectly informed. Topic-specific biases, while probably not as massive as previously assumed, are nevertheless likely to be widespread. The model can easily be adapted to account for topic-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 topic-specific biases can help account for the small residual errors that are not explained by domain-general, rational rescaling. Media bias, for example, 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.

In past research on how perceptions affect political judgments and decisions, most researchers have interpreted raw estimation error as a direct measure of individual misinformation or bias (but see Kuklinski, et al., 2000 and Pasek, Sood, & Krosnick, 2015). However, from the psychophysical perspective for which we are advocating here, political

expressions of beliefs will be driven, not just by perceptions themselves, but by the outcome of some process that transforms perceptions into responses. As a result, citizens may internally hold unbiased perceptions about demographic groups, but the unavoidable process of transforming those perceptions into responses will introduce systematic distortions, whether citizens are estimating proportions for a survey or using their perceptions to decide how vote. Making sense of these different transformations, which can differ by context or task, will improve our understanding of individual-level political behavior and help develop interventions aimed at correcting perceptions and civic behavior. For instance, people who feel more threatened by minorities also provide higher estimates of the size of minority populations (Nadeau et al., 1993). But such differences in overestimation may reflect two very different phenomena among respondents who feel more threatened: greater anti-minority bias (the preferred explanation of many political scientists) or increased uncertainty about the size of minority populations (which, on our account, will also lead to overestimations). Without a proper analysis that is grounded in the psychology of quantity estimation, it is impossible to distinguish these different sources of error.

Identifying the true source of errors in demographic estimations is important both practically and theoretically. Most obviously, if a person does not know at all how to transform a perceived quantity into a reported percentage, then 'informing' them of the true immigration rate may be essentially worthless. Alternatively, if people are massively misinformed but give reasonable estimates due to uncertainty-related psychophysical transformations, then incorrect perceptions may be left unaddressed by interventions that specifically target issues that are reliably misestimated. These possibilities may contribute to previous findings that supplying relevant information, like the rate of unemployment and poverty in the U.S., has limited utility in

shifting estimates and political opinions (Lawrence & Sides 2014). Similar psychophysical analyses may play a helpful supporting role in aiding interpretation of different patterns of overestimation across states, counties, and individuals (Wong, 2007), but we leave this to future research.

While past work on political ignorance has largely 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 a common value (e.g., Hastie, Tibshirani, & Friedman, 2009). From a statistical perspective, this can be 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 the principle of uncertainty-based rescaling. 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.

There are some limitations to our analytic framework. 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 model used here is the best among those psychophysical models. Our intention 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 topic-specific, ad-hoc explanations. Second, here we analyzed group-level data, but the model was developed to account for individual behavior (Gonzalez & Wu, 1999). This is problematic because the 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 Ipsos-Mori dataset 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 fairly 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 rational rescaling of log-odds (Marghetis et al, in prep).

Two outstanding questions remain unanswered: First, do individual differences in 'uncertainty-based rescaling' of demographic proportions reflect differences in individuals' uncertainty? On our account, the more uncertain an individual is about a true value, the more they should rely on the prior distribution of values typical to that category (see also Stewart et al, 2006). 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, 2016; Fox & Rottenstreich, 2003, Martins, 2006). Second, how do an individual's demographic estimates relate to their political behavior? That is, when people engage in political behavior, such as voting or advocating for a policy, do they rely primarily on their explicit estimates — which combine perceptions with prior expectations — or do they behave on the

basis of perceptions alone? In at least some decision-making contexts, rescaled probability estimates have been found to predict subsequent evaluations and behaviors (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, perceptions may be more potent influences on behavior.

These results have important implications for policy makers, the press, and political scientists. Overestimation of minority populations has often been taken as evidence of personal or topic-specific biases—most commonly media bias, phobic innumeracy, or social contact biases. Our results underline the principle that overestimation is a natural consequence of making estimates about values that are smaller than the expected or typical value (e.g., the proportion of the U.S. population that is Asian, which is smaller on average than most demographic groups). In order to conclude that there is a truly topic-specific bias, individuals' demographic estimates need to exceed what we should predict on the basis of rational 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 (Druckman et al. 2009). Similarly, psychologists are unlikely to understand the full ramifications of 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 important theoretical and empirical questions.

5   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.

6   Acknowledgments

7   References

Alba, R., Rumbaut, R. G., & Marotz, K. (2005). A distorted nation: Perceptions of racial/ethnic group sizes and attitudes toward immigrants and other minorities. *Social Forces*, 901-919.

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.

Berinsky, A. J. (2015). Rumors and health care reform: experiments in political misinformation.

British Journal of Political Science, 1-22.

Birnbaum, M. H., & McIntosh, W. R. (1996). Violations of branch independence in choices

between gambles. Organizational Behavior and Human Decision Processes, 67, 91–110.

Box, G. E., & Tiao, G. C. (1973). Bayesian Inference in Statistical Analysis. Addison Wesley,

Reading, Massachussetts.

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. (2016). 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)

Druckman, James N., James H. Kuklinski, and Lee Sigelman. 2009. "The Unmet Potential of Interdisciplinary Research: Political Psychological Approaches to Voting and Public Opinion." *Political Behavior*, Vol. 31, No. 3: 485-510.

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.

Gallagher, C. A. (2003). Miscounting race: Explaining Whites' misperceptions of racial group size. *Sociological Perspectives, 46*(3), 381-396.

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.

Gescheider, G. A. (1976). Psychophysics: Method and theory.

Gilens, D. (2001). Political ignorance and collective policy preferences. *American Political Science Review, 95,* 379-396.

Goldstein, W. M., & Einhorn, H. J. (1987). Expression theory and the preference reversal phenomena. *Psychological Review*, 94, 236–254.

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.

Herda, D. (2010). How many immigrants? Foreign-born population innumeracy in Europe. *Public opinion quarterly*, *74*(4), 674-695.

Herda, D. (2013). Too Many Immigrants? Examining Alternative Forms of Immigrant Population Innumeracy. *Sociological Perspectives*, *56*(2), 213-240.

Herda, D. (2015). Beyond innumeracy: heuristic decision-making and qualitative misperceptions about immigrants in Finland. *Ethnic and Racial Studies*, *38*(9), 1627-1645.

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

Karmarkar, U. (1979). Subjectively weighted utility and the Allais paradox. *Organizational
    Behavior and Human Performance*, 24, 67–72.

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. (2016). Categories of Large Numbers in Line
    Estimation. *Cognitive Science,* early online publication. 10.1111/cogs.12342

Lattimore, P. K., Baker, J. R., & Witte, A. D. (1992). The influence of probability on risky
    choice: A parametric examination. *Journal of Economic Behavior and Organization,* 17,
    377–400.

Lawrence, E. D., & Sides, J. (2014). The consequences of political innumeracy. *Research &
    Politics*, *1(2),* 2053168014545414.

Lee, M. D., & Danileiko, I. (2014). Using cognitive models to combine probability
    estimates. *Judgment and Decision Making*, *9*(3), 259.

Lodge, M., & Taber, C. S. (2013). *The Rationalizing Voter*. (New York: Cambridge University
    Press)

Lundmark, S., & Kokkonen, A. (2014). Measuring Immigrant Populations: Subjective versus
    Objective Assessments. *APSA 2014*: Annual Meeting Paper.

Lupia, A., & McCubbins, M. D. (1998). *The democratic dilemma: Can citizens learn what they need to know?*. Cambridge University Press.

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.

Mosteller, Frederick (1977). Data analysis and regression : a second course in statistics. Reading, Mass: Addison-Wesley Pub. Co. ISBN 978-0201048544.

Nadeau, R., R. G. Niemi, and J. Levine. 1993. "Innumeracy About Minority Populations." *Public Opinion Quarterly 57*(3), 332-47.

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)

Nyhan, B., & Reifler, J. (2010). When corrections fail: The persistence of political misperceptions. Political Behavior, 32(2), 303-330.

Pasek, J., Sood, G., & Krosnick, J. A. (2015). Misinformed about the Affordable Care Act? Leveraging certainty to assess the prevalence of misperceptions. *Journal of Communication*, *65*(4), 660-673.

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.

Petzschner, F (2012). Magnitude estimation in humans: A Bayesian Approach to Characteristic Behavior in Path Integration. Doctoral Dissertation for the Graduate School of Systemic Neurosciences at the. Ludwig-Maximilians-Universität München

Petzschner, F. H., Glasauer, S., & Stephan, K. E. (2015). A Bayesian perspective on magnitude estimation. Trends in cognitive sciences, 19(5), 285-293.

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., Kilpatric, 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.

Stevens, S. S. (1975). *Psychophysics*. Transaction Publishers.

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology, 53*, 1-26.

Tversky, A., & Fox, C. R. (1995). Weighing risk and uncertainty. *Psychological Review*, 102, 269–283.

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

8   Appendix

The European Social Survey (ESS) was conducted in 2002-2003 across 20 European countries. The mean response rate across all countries was 60%. While all interviews were conducted in-person, the sampling design varied across countries (for details on sampling design and response rates by country, see http://www.europeansocialsurvey.org). The analysis presented here does not include estimates from Israeli and Slovenian respondents, as benchmark immigration data are not available for these countries.

The Ipsos MORI Perils of Perception was conducted in 14 countries across Europe, North America, and Asia between August 12th-14th, 2014. Respondents from all countries were members of the Ipsos Online Panel, a non-probability sample of internet-users. The data are weighted by Ipsos Mori to appear representative of each country's population. Respondents were asked to estimate how many, out of every 100 people, in their country had certain characteristics, such as being Muslim, unemployed, or elderly.