

CHOOSING YOUR POND: LOCATION CHOICES AND RELATIVE INCOME

Nicolas L. Bottan and Ricardo Perez-Truglia*

Abstract—Do individuals care about their relative income? While this is a long-standing hypothesis, revealed-preference evidence remains elusive. We provide a unique test by studying residential choices: individuals often must choose between places with different income distributions, and as a result they “choose” their relative income. We conducted a field experiment with 1,080 senior medical students who participated in the National Resident Matching Program. We estimate their preferences by combining choice data, survey data on perceptions, and information-provision experiments. The evidence suggests that individuals care about their relative income and that these preferences differ across single and nonsingle individuals.

I. Introduction

DO individuals care about their relative income? Economists have been pondering this possibility since the seminal work of Adam Smith (Luttmer, 2005). Consistent with the hypothesis that they do indeed care, a large literature shows that subjective well-being increases with relative income (Senik, 2004; Luttmer, 2005; Ferrer-i-Carbonell, 2005). The revealed-preference evidence, however, remains elusive. In this paper, we provide unique revealed-preference evidence based on the choice of reference groups.

One infrequent but important choice individuals must make is where to live. By choosing where to live, individuals also choose their reference group, the other individuals with whom they will compare their incomes. If individuals choose to move to a poor area, they will be relatively richer than their peers. If they choose to move to a rich area, they will be relatively less rich than their peers. In this paper, we estimate preferences for relative income by studying the location choices of 1,080 senior medical students from the United States.

The identification of preferences for relative income, or any other city amenity, faces a number of data and identification challenges.¹ We developed a novel methodology

Received for publication January 7, 2020. Revision accepted for publication October 1, 2020. Editor: Shachar Kariv.

*Bottan: Cornell University; Perez-Truglia: University of California, Berkeley.

We are thankful for comments from several colleagues and seminar discussants. We thank UCLA-Anderson and its Behavioral Lab for providing funding for the experiment; and thank the support of the Robert Ferber Dissertation Award. This project was reviewed and approved in advance by the Institutional Review Board at University of California Los Angeles (IRB 16-001968; 17-001449). AEA RCT Registry 0004203.

A supplemental appendix is available online at https://doi.org/10.1162/rest_a_00991.

¹For example, one challenge consists of observing not only the city and amenities that the individual chose but the cities and amenities that the individual could have chosen instead. Some data sets, such as the National Survey of Families and Households (Luttmer, 2005) and the U.S. Postal Service’s National Change of Address database (Perez-Truglia, 2018), identify individuals moving from one city to another, but do not identify the other cities that the individual could have chosen or incomes they could have earned in those other locations. Another challenge is that we need sources of exogenous variation to identify the amenity preferences. In other words, preferences for a given amenity could be spuriously driven by preferences for another amenity that is not included in the regression.

to address those challenges that combines choice data, survey data and an information-provision experiment. While our methodology can be applied to any sample of individuals who are choosing where to live, we implemented it in the specific context of the National Resident Matching Program (NRMP).

The NRMP uses an algorithm to pair graduating medical students with postgraduate hospital residency programs based on rankings submitted by both students and hospitals. When submitting their program rankings, the students are choosing the city where they will live for the duration of the residency (typically around five years). Several features of the NRMP make it desirable for this type of revealed-preference analysis (Benjamin et al., 2014). First, the NRMP makes it possible to identify and specify the entire choice set these individuals face. Second, because most students are aware of the incentive-compatible matching algorithm used by the NRMP, it is possible to infer preferences directly, without the need for estimating models that rely on additional assumptions. Third, this is a high-stakes choice, to which participants devote ample time and attention. A fourth and convenient aspect of our setting is that the graduating students choose between programs that offer almost identical nominal incomes but in cities with largely different income distributions and costs of living. As a result, our subjects face substantial trade-offs between relative income and cost of living. Finally, this context provides a stringent test for estimating preferences for relative income; aspects such as residency quality and prestige are likely the most predominant in medical student’s decisions.

We conducted the experiment with 1,080 senior medical students who participated in the 2017 residency match. We conducted a baseline survey roughly about one month before these participants had to submit their rank-order preferences. We asked participants about their two favorite residency programs. Although students can rank several programs, the ranking of their top two is, in practice, the most important aspect of their decision. We elicited perceptions about their expected position in the city-wide distribution of individual earnings (i.e., their relative income). We also elicited perceptions about another amenity that is thought to be closely related: their expected cost of living. We then elicited each subject’s expected rank submission. Using these data on perceptions and choices, we can estimate how differences in the cities’ relative incomes and costs of living affect location choices.

These estimates of location preferences may be subject to omitted variable biases. We use two strategies to address this identification challenge. The first relies on controlling for a host of alternative factors such as the reputation of the residency programs and all the typical city amenities that have been studied in the urban economics literature. The second strategy relies on an information-provision experiment.

Immediately after eliciting perceptions about cost of living and earnings rankings, we provided all individuals with statistics about attributes of the cities where their programs are located. We randomized the value of this feedback in a non-deceptive way by randomizing the data source used to compute these statistics. For instance, students who expected to earn \$54,000 at a residency in Champaign-Urbana, Illinois, received one of two messages: their earnings rank would be 55.1% according to data from the Current Population Survey or 60.3% according to data from the American Community Survey. We then elicited their perceptions again after providing this feedback (i.e., their posterior beliefs). This source-randomization experiment creates exogenous variation in posterior beliefs. Our research design exploits that exogenous variation to estimate the causal effects of the perceived attributes on choice by means of an instrumental variables model.²

Our motivation for studying relative income lies in models of social status and consumption aspirations. According to these models, the income of neighbors can impose a negative externality, thus giving individuals an incentive to locate in less affluent areas (Luttmer, 2005). However, individuals may care about the income of their neighbors for other reasons. Moreover, some individuals may see the income of neighbors as a positive externality, thus giving them an incentive to move into more affluent areas. For example, certain individuals may expect to have better dating prospects in more affluent areas. Alternatively, individuals may expect to benefit from public goods provided by richer neighbors.

To maximize statistical precision, our baseline estimates exploit all the variation in expectations, experimental and nonexperimental. These estimates suggest that, holding cost of living constant, the average individual would prefer to live in a city where her relative income would be higher. More precisely, a 1 percentage point increase in earnings rank on average increases the probability that a program will be chosen by 0.186 percentage points (and implying a behavioral elasticity of 0.186).

We use cost of living as a benchmark for relative income. Our baseline estimates suggest that a 1 percentage point decrease in cost of living, on average, increases the probability that a program will be chosen by 0.201 percentage points (and implying a behavioral elasticity of 0.201). These estimates imply that the average individual would be willing to accept an increase of 0.925 percentage points in the cost of living in exchange for a 1 percentage point increase in earnings rank. This trade-off suggests that the preferences for relative income are economically significant.

The importance of financial aspects such as relative income and cost of living is consistent with survey data indicating that students chose medical school because of the financial

rewards (Daniel & O'Brien, 2008). Additionally, these findings are consistent with survey data indicating that, when choosing among programs, applicants care a lot about the geographic location of the programs (Sledge, Leaf, & Sacks, 1989). Indeed, when asked to mention the most important factors they consider in their decision making, 52% of surveyed applicants explicitly mention cost of living (National Resident Matching Program, 2015). Although the preferences for relative income and cost of living are statistically and economically significant, they do not imply that they are the main concern for medical students. Indeed, using perceptions about other program characteristics, we find that doctors care substantially more about prestige and career prospects, which is consistent with the top factors that medical students mention in surveys (NRMP, 2015).

We show that these average preferences mask meaningful heterogeneity by relationship status. While nonsingles prefer being a big fish in a small pond, single individuals prefer being a small fish in a large pond. This difference in preferences is large in magnitude and highly statistically significant (p -value < 0.001). This heterogeneity is consistent with evidence from the happiness literature: Luttmer (2005) finds that the positive effects of relative income on happiness are driven entirely by nonsingles. However, when it comes to preferences for cost of living, individuals have similar preferences regardless of whether they are single or nonsingle.

One important caveat with the baseline estimates is that relative income may be picking up the effects of omitted variables. We attempt to address these concerns with two strategies. The first shows that the estimates remain similar after controlling for a host of other characteristics of the residency programs and the cities where they are located. In the second strategy, we focus on the variation in beliefs generated by the information-provision experiment. The estimates are less precisely estimated because we are focusing on a minority of the variation in beliefs, but the main results remain robust.

However, even the experimental estimates are subject to interpretation: individuals may react to the information about relative income not because they care about relative income per se but because it serves as a signal for other unobservable attributes of a city. While we cannot totally rule it out, we provide two pieces of evidence against this interpretation. We show that individuals are not using the information on earnings rank to update their expectations about cost of living, in both the short term (the baseline survey) and a longer horizon (the follow-up survey). Second, we conducted an auxiliary experiment with a sample of respondents recruited through Amazon Mechanical Turk. In this complementary experiment, we provide feedback about earnings rank and cost of living, but we measure posterior beliefs not only about earnings rank and cost of living but also about other characteristics of the city that the individuals may be learning about based on the earnings rank information (e.g., public goods, crime rates). We show that the results are, again, unchanged when we control for this host of posterior beliefs.

²Our design differs from the traditional approach used in information experiments that consists of using a single source of feedback, but randomly showing feedback to half of the subjects. Instead, we provide information to all subjects and randomly nudge some subjects to perceive that they will be richer through the source randomization.

Our favorite explanation for the preference for less affluent ponds, though it is by no means the only explanation, is that individuals see their neighbor's income as a negative externality. Moving to a more affluent pond, for instance, may result in a loss of social status (Luttmer, 2005; Bénabou & Tirole, 2006; Ray & Robson, 2012), a loss of self-esteem (Festinger, 1954), higher consumption aspirations (Frank, 1985b), and poorer outcomes in social interactions (Doob & Gross, 1968; Fennis, 2008; Nelissen & Meijers, 2011). Regarding the heterogeneity by relationship status, our preferred interpretation is that single individuals expect more positive externalities from affluent neighbors such as enjoying public goods and having better dating prospects.

Our paper is related to several strands of literature. First, it relates to literature on the importance of relative income for subjective well-being. Since the seminal contribution of Easterlin (1974), several studies have argued that subjective well-being depends on relative, rather than absolute, income (Clark & Oswald, 1996; Senik, 2004; Luttmer, 2005; Ferrer-i Carbonell, 2005; Perez-Truglia, 2016).³ Luttmer (2005), for instance, uses data from the United States to show that, holding own income constant, happiness increases with relative income in the area of residence. We contribute to this literature in two ways. Our main contribution is to develop a method to test the relative income hypothesis by using revealed-preference data instead of happiness data. Additionally, even if we take for granted that the income of neighbors imposes a negative externality on happiness, it is unclear whether individuals anticipate these externalities when deciding where to live. Our results provide suggestive evidence that some individuals may anticipate these externalities, at least partially.

This paper is related to a strand of literature using laboratory experiments to study positional externalities. Some of these studies use surveys that let subjects choose between pairs of hypothetical scenarios that encompass trade-offs between absolute and relative standings. These studies find that individuals are often concerned about their relative standing (Solnick & Hemenway, 1998; Johansson-Stenman, Carlsson, & Daruvala, 2002; Yamada & Sato, 2016; Clark, Senik, & Yamada, 2017). In a similar spirit, some laboratory experiments show that relative standing affects behavior in laboratory games such as choices under risk (Kuziemko et al., 2014). We contribute to this literature by estimating the trade-offs between absolute and relative income in a real-world, high-stakes context.

This study is also related to a strand of literature on social status (Frank, 1985a; Heffetz, 2011). For example, Bursztyn et al. (2017) conducted a field experiment to show that individuals use platinum cards to signal their social status. This

³These studies often use a slightly different specification: holding own income constant, well-being decreases with the average income in the reference group. It must be noted that some studies find the opposite effect (Senik, 2004) or mixed evidence (Clark, Westergård-Nielsen, & Kristensen, 2009). For an extensive review of the literature, see Tideman, Frijters, and Shields (2008).

literature, however, takes reference groups as exogenous and fixed. As a result, there is little direct evidence on how individuals select into reference groups. If individuals truly care about their relative income, when given the choice, they should choose to live in places where their relative income is expected to be higher (Frank, 1985a). Since this is a basic prediction of models of social status, it is notable that this prediction of status has not been tested before. Our contribution is to fill this gap in the literature.

Our paper is additionally related to a large literature measuring preferences for city amenities. Starting with the model pioneered by Rosen (1979) and Roback (1982), economists have been using structural models to estimate the value of amenities (Albouy, 2008, 2016). However, there is a growing demand for credible sources of identification. A few studies have exploited quasi-experimental sources of identification to estimate the value of specific amenities. Black (1999), for example, exploits a geographic discontinuity in school district borders in order to estimate preferences for school quality (see also Bayer, Ferreira, & McMillan, 2007). However, due to the limited availability of natural experiments, the quasi-experimental evidence is limited to a small number of amenities and contexts. The methodology developed in this paper can be used to estimate preferences for other types of amenities and contexts. This methodology can provide credible identification of preferences without the need for natural experiments.

Finally, our study is also related to a strand of literature finding that individuals substantially misperceive their relative incomes (Cruces, Perez-Truglia, & Tetaz, 2013; Karadja, Mollerstrom, & Seim, 2017).⁴ This literature shows that correcting these misperceptions has significant effects on stated preferences for redistribution. Yet there is no evidence that these misperceptions have a significant effect on behavior. We contribute to this literature by filling that gap.

The rest of the paper proceeds as follows. Section II describes the survey design. Section III presents the econometric model. Section IV presents implementation details and descriptive statistics. Section V discusses the distribution of perceptions and learning. Section VI presents results on location preferences.

II. Survey Design

A. *Timing of the Surveys*

Inspired by Benjamin et al. (2014), we study the decisions of medical students participating in the NRMP.⁵ After graduating from medical school, students have to complete a

⁴Misperceptions have been reported on other aspects of society too, for example, about the degree of income inequality (Hauser & Norton, 2017).

⁵Benjamin et al. (2014) conducted a survey of medical students after the students had submitted their rankings to the NRMP. The survey measured these rankings as well as the perceived rank of many aspects of the programs, such as life satisfaction, happiness, and sense of control. They measure and compare the preferences inferred from rank choices to those inferred from subjective well-being. We follow this survey collection method closely,

residency to become a *medicinae doctor* (MD). A residency usually lasts from three to seven years, after which individuals may obtain their medical license. During the fall semester of 2016, fourth-year medical school students started their participation in the residency match by submitting applications to residency programs. Later in the semester, they were interviewed and flown out by some of the programs they applied to.⁶ After all interviews were completed, the students spent almost two months deciding how to rank their favorite programs. During this time, deciding on the rank-order preference is the students' top priority: applicants claim to collect a lot of information to aid their decision, such as characteristics of the residency programs and characteristics of the cities where the programs are located, and they may even visit the cities again. In the 2017 match, the submission window for rank-order lists opened on January 15 and closed on February 22. We conducted a baseline survey weeks before the submission deadline, which we describe first, and a follow-up survey right after the submission window closed.

B. General Structure of Baseline Survey

The baseline survey starts and ends with some background questions, such as the subject's medical school and marital status (see appendix C.1 for the full questionnaire of the baseline survey). The core of the survey comprises the following group b of questions, in the order listed below:

1. *Choice set*: Elicit the names of the two favorite programs that the individual was considering for his or her order rank submission.
2. *Prior beliefs*: Elicit perceptions about the earnings rank and the cost of living in the cities where these two programs are located.
3. *Feedback*: Provide subjects with feedback related to their perceptions.
4. *Posterior beliefs*: Reelicit perceptions about the earnings rank and the cost of living.
5. *Rank choice*: Elicit the individual's expected rank submission (between the two programs).

The following sections provide details about each of these modules.

Choice set. The survey asks individuals to list their top two preferred programs, in no particular order, from a user-friendly list of all the available programs organized by state and metropolitan area. While in theory students could list any of the residency programs in the country, in practice they tend to focus on the programs that they think may be interested in them. For example, if one student has a flyout at

a particular program, that is a strong reason to list it in the ranking.⁷ Medical students are mostly decided on a specialty when submitting their rankings (indeed, only 1.6% of our sample indicate programs in different specialties). We limited the survey to two programs because otherwise it would have been too cognitively demanding. Most participants expect to be matched to one of their top-two choices: similar to previous years, 50.9% of the participants in the 2017 match were assigned to their first choice and 16.6% were assigned to their second choice. We concentrated on the participants' top two programs rather than a random pair of programs because this happens to be the part of the decision with the highest stakes and to which individuals were paying the most attention. In any case, our focus on the top two choices does not challenge the validity of our estimates: the research design would be valid with any pair (or group) of options, not only the top two.⁸

Perceptions about earnings rank and cost of living. One important feature of the residency match process is that salaries are relatively homogeneous across the different programs, even across specialties.⁹ Indeed, each program offers the same salary to all its candidates (and that salary is often publicly available on the program's website). Despite the homogeneity in nominal incomes, there is large heterogeneity in costs of living and earnings distributions in the cities where the programs are located. When designing the survey, we were constrained to using metropolitan areas rather than other geographical levels of aggregation (e.g., commuting zones) because the sources of data on cost of living are not collected at a finer level than the metro area.

We asked two questions about earnings rank (one for each metro area) and two questions about the cost of living (one for each metro area). We provided an introduction for the question about earnings rank: "Now we want to ask you about your expected earnings rank. This rank is defined as the share of the working individuals of a city who earn less than you. You probably noticed that the distribution of earnings is different across different cities. As a result, with the same earnings, you may be relatively rich in some cities but relatively poor in other cities." After this introduction, we asked the following question for each city: "Imagine that you chose to

⁷For more details on the supply and demand in this labor market, see, for example, Roth and Peranson (1999) and Agarwal (2015).

⁸When individuals were listing the second program, we required respondents to make a selection from a different metro area because otherwise no differences would be present in relative income and cost of living across choices. Our survey data indicate that no more than 6% of individuals tried to select the same metro area. For those subjects, the comparison was between two of their top programs but not necessarily the top two. In appendix A.9, we show that the results are robust if we exclude those individuals from the sample.

⁹Even though there are no large income differences in residency salaries, there can be large differences in postresidency salaries, especially across specialties. However, this is not a concern because medical students are choosing programs within a specialty. Furthermore, the differences in students' perceptions of postresidency salaries among different programs within specialty are small (Benjamin et al., 2014).

but we change the survey itself to test a different hypothesis. Our survey does differ in some important aspects: we collected our baseline survey before subjects submit their rank choices to the NRMP and embedded an information-provision experiment into the survey.

⁶In 2015, the median number of applications submitted was thirty and the median number of interviews sixteen (NRMP, 2015).

work in [metro name]. With your individual annual earnings of \$[salary], you would be richer than what percentage of [metro name]'s individual earners?" Respondents could select their answer from a drop-down menu that ranges from "Richer than 1% of individual earners" to "Richer than 100% of individual earners," in 1% increments.¹⁰

For the cost-of-living question, we stated: "You probably noticed that the average prices of goods and services are different across different cities. As a result, with the same income, you would be able to buy more things in some cities and less in other cities." After this introduction, we asked individuals how much more or less expensive each metro area was relative to the U.S. average. To make answering the question easier, we split it into two questions. The first question was: "Imagine that you chose to work in the [metro name] metro area. Would you expect your cost of living in this city to be cheaper or more expensive than the U.S. average?" The respondents could choose either "cheaper" or "more expensive." The second part of the question was: "How much [cheaper/more expensive] is the [metro name] metro area than the U.S. average?" Respondents could answer this second question with a drop-down menu ranging from 0% to 50%, in 1 percentage point increments.

We focus on this definition of reference group because it is the most widely used approach in the related literature: Luttmer (2005), for example, studies how the happiness of an individual is affected by the income of her neighbors.¹¹ In practice, individuals may care about their ranking in finer reference groups; for example, they may care disproportionately about their relative standing with respect to neighbors in the same age cohort rather than caring about all neighbors equally. However, this source of measurement error is not a major source of concern, to the extent that it can only introduce attenuation bias.

Information-provision experiment. One limitation with using perceptions is the potential for omitted-variable bias. For instance, conditional on income and perceptions about cost of living, perceptions about relative income may happen to be correlated with perceptions about other characteristics of the area, such as the crime rate, amenities, public goods, and so forth. To address this concern, we generate exogenous variation in the perceptions about earnings rank and cost of living by embedding an information-provision experiment in the survey.

Immediately after respondents provided their prior beliefs on both measures, they were shown two messages: one page with statistics about the earnings rank in the two cities being considered and another page with statistics about the cost of

living in each of the two cities. The following message is a sample of the feedback page for earnings rank: "With your individual annual earnings of \$54,000, you would be richer than 57.9% of Los Angeles-Long Beach-Anaheim, CA's population. With your individual annual earnings of \$54,000, you would be richer than 60.3% of the population of Champaign-Urbana, IL's population." The following message is a sample of the feedback page for cost of living: "Los Angeles-Long Beach-Anaheim, CA metro area is 17.0% more expensive than the U.S. average. The Champaign-Urbana, IL metro area is 6.6% cheaper than the U.S. average." In both of these feedback pages, individuals were asked to take a moment to review the information carefully and were alerted that the information was going to be shown only once. We did not allow respondents to continue to the next page until at least 10 seconds had elapsed.¹²

After individuals finished reviewing the feedback, we reelicited their perceptions about these attributes for each city, which we denote as the posterior beliefs. Given that our feedback entailed many figures for participants to remember and process, we wanted to make it easier for individuals to compare the options. Therefore, after eliciting respondents' posterior beliefs, we gave them a third page of feedback based on their posterior beliefs. The following is a sample of that feedback page: "We understand this is a lot of information to process, so we will help you make the comparison simpler. According to your final answers about incomes, cost of living and earnings rank: If you chose to live in Los Angeles-Long Beach-Anaheim, CA, you would be able to afford 19.7% less than if you chose to live in Champaign-Urbana, IL. If you chose to live in Los Angeles-Long Beach-Anaheim, CA, your earnings rank would be 3.3% lower than if you chose to live in Champaign-Urbana, IL."¹³

We computed the statistics shown to the subjects using two alternative data sources, and we cross-randomized which of the two sources were shown to each individual. The sources were randomized between individuals; that is, we used the same cost-of-living source for the two cities being considered by each individual and the same earnings data source for the two cities. As a result, individuals were randomly assigned to one of four treatment groups. For the earnings rank feedback, the two sources used were the American Community Survey (ACS) and the Current Population Survey (CPS), both conducted by the U.S. Census Bureau. For cost-of-living estimates, the two sources used were the Regional Price Parity (RPP) data by the Bureau of Economic Analysis and the

¹²The median time spent on the feedback page was 18.5 seconds.

¹³The difference in cost of living was calculated as $100 \times \left(\frac{w_1 \text{COL}_2}{w_2 \text{COL}_1} - 1 \right)$, where w_i is the nominal wage for city i and COL_i is their posterior belief about cost of living (from 50 to 150). The difference in earnings rank was calculated as $100 \times \left(\frac{\text{ER}_1}{\text{ER}_2} - 1 \right)$, where ER_i is the posterior belief about earnings rank in city i . As with the other feedback pages, 10 seconds had to elapse before respondents could move to the next page. The median duration on the post-feedback page was 19.5 seconds.

¹⁰We decided not to incentivize any of the elicitation—for details, see the discussion in appendix A.6.

¹¹Moreover, this geographic definition of references group is used more generally in the literature of social interactions more generally; for example, Perez-Truglia (2018) and Perez-Truglia and Cruces (2017) study how an individual's political participation is affected by the participation of her neighbors.

Cost of Living Index (COLI) data compiled by the Council for Community and Economic Research.¹⁴

This source randomization created a substantial amount of exogenous variation in signals. For instance, the correlation of the pairwise difference in earnings rank shown to the respondents versus the corresponding pairwise difference from the alternative source is 0.649; the corresponding correlation for the cost of living is 0.656. These differences across sources arise from a combination of several factors, most notably sampling variation and data definitions. For instance, the earnings rank data are subject to sampling variation because the estimates are based on a limited number of survey respondents, and cost-of-living data are subject to sampling variation because they track the prices of a limited number of goods and services. The variation in definitions arises because the earnings rank measures are based on surveys with significant differences in the survey method, and the phrasing of the questions used to elicit total annual earnings, and different cost-of-living indices, give different weights to expenditure categories.

For transparency and to ensure the validity of the information, the individuals were debriefed in the feedback messages on the name of the source of the information that they received. We would not expect the source name to have an effect in and of itself, given that the individuals did not have expertise on the data. Indeed, we find that the reaction of individuals to the information was orthogonal to the name of the information source.¹⁵

Rank submission choices. The survey asked respondents to indicate which program they expected to rank higher when submitting to the NRMP: “As of this moment: of the two programs discussed so far, which one would you expect to rank higher for the NRMP?” Individuals could indicate their ranking on a 6-point scale ranging from “Very likely [Program 1] (in [Metro 1])” on one side to “Very likely [Program 2] (in [Metro 2])” on the other. In the baseline results we look at the binary choice of whether they expect to rank program 1 over program 2 because a comparison with the ex post submission choices is more straightforward. Nevertheless, results are similar when using the full likelihood scale (results reported in appendix A.17).

It is plausible to assume that the rank choices provide a great proxy for the individuals’ true preferences. First, the individuals have strong incentives to take the submission seriously, as whatever they submit is quite consequential: almost all NRMP participants receive a match,¹⁶ and backing out from a match entails serious sanctions.¹⁷ Second, the al-

gorithm used by NRMP was designed by Roth and Peranson (1999) to be 100% resistant to attempts of “strategic behavior,” meaning that it is a weakly dominant strategy for students to submit their true preferences. However, even if the mechanism is strategy-proof in theory, that does not imply that the participants are reporting the truth in practice; for example, they may misunderstand the rules of the mechanism. One nice feature of the setting is that the NRMP provides training to the students, making it explicit that it is in their best interest to report preferences truthfully. Moreover, there is some evidence indicating that the training works; for example, survey data from Benjamin et al. (2014) indicate that only around 5% of participants attempt to misreport their true preferences with a strategic motive (see also NRMP, 2015, and Rees-Jones, 2018).¹⁸ In other words, while the rankings submitted to the algorithm may not be exactly their true preference, they are likely a close approximation.

C. Follow-Up Survey

Shortly after the NRMP rank submission window closed, we conducted a follow-up survey with the subjects who responded to the baseline survey. Appendix C.2 shows the full questionnaire of the follow-up survey. Most important, at the beginning of the survey, we collected data on the final rank order submitted to the NRMP. We also took the opportunity to ask individuals for some additional information. We elicited the perceptions about earnings rank and cost of living, which allows us to measure the persistence of the information learned in the information-provision experiment. Also, we measured additional characteristics of the subjects that we were not able to measure in the baseline survey due to space and time constraints, such as the places where the individuals grew up.

III. Econometric Model

A. Baseline Model

In this baseline model, we exploit all the variation in perceptions of earnings rank and cost of living, which includes the experimental variation induced by our information provision, as well as the remaining nonexperimental variation.

Let i index subjects and $j \in \{1, 2\}$ denote the two programs being considered by the subject. We define $ER_j^{i, \text{posterior}}$ and $COL_j^{i, \text{posterior}}$ as the posterior beliefs for earnings rank and cost of living for program j in the baseline survey.

lifetime bar from participation in future NRMP matches, and from one year to a lifetime identification in the matching system as a match violator (<http://www.nrmp.org/policies/the-match-commitment/>). Additionally, the NRMP has established rules prohibiting programs from contacting candidates to ask or coordinate their rank orders.

¹⁸There is a growing literature on how strategy-proof mechanisms work in practice. See, for example, Li (2017) and Rees-Jones and Skowronek (2017).

¹⁴For more details, see appendix A.1.

¹⁵Results reported in appendix A.7.

¹⁶For instance, 95% of the 27,048 U.S. graduating medical students received a successful match in 2017.

¹⁷For example, applicants with confirmed violations of NRMP policies are subject to a one-year bar from accepting or starting a position in any program sponsored by a match-participating institution, from one year to a

Let $ER_{1,2}^{i,posterior} = ER_1^{i,posterior} - ER_2^{i,posterior}$ be the perceived difference in earnings rank between the two programs. Similarly, let $COL_{1,2}^{i,posterior} = COL_1^{i,posterior} - COL_2^{i,posterior}$ be the perceived difference in cost of living between the two programs. Let $Program_1 \succ_i Program_2$ denote that individual ranks program 1 over program 2, and let $I(\cdot)$ be an indicator function. The regression specification is

$$I(Program_1 \succ_i Program_2) = I(\beta^{ER} \times ER_{1,2}^{i,posterior} + \beta^{COL} \times COL_{1,2}^{i,posterior} + \theta X^i + \varepsilon_i \geq 0), \quad (1)$$

where X^i is a vector of control variables and θ is the corresponding vector of coefficients. We always include a constant and the log-difference of nominal residency wages as control variables. In the baseline specification, we include an additional set of controls consisting of pairwise differences in some residency and location characteristics: residency program rank (from Doximity), quality of life inferred from compensating differentials (Albouy, 2016), population size, population density, share of African American residents, share of Democrat residents, and share of urban population.¹⁹ In any case, we present results with alternative sets of control variables.

In the baseline specification, we estimate a probit model, which implies that the error term (ε_i) is normally distributed. As is typical in discrete-choice models, using a probit model is convenient in the sense that the ratio between parameters can be readily interpreted as marginal rates of substitution. However, this specification choice is irrelevant in practice: the results are very similar if we use alternative models such as binary logit, ordered probit, or linear probability (results reported in appendix A.16).

The two key parameters of interest are β^{ER} and β^{COL} . The parameter β^{ER} measures preferences for relative income during the residency. Depending on the mechanism at play, we may expect β^{ER} to be positive or negative. On the one hand, if individuals see their neighbor's income as a negative externality, as in the models of social status, then we would expect $\beta^{ER} > 0$ (i.e., individuals want to choose less affluent ponds). On the other hand, if individuals see the neighbor's income as a positive externality, then we would expect $\beta^{ER} < 0$ (i.e., individuals want to choose more affluent ponds). As our benchmark, the parameter β^{COL} measures preferences for purchasing power during the duration of the residency. The hypothesis is that $\beta^{COL} < 0$: individuals prefer to live in places with a lower cost of living.

The duration of a residency depends on the specialty: it lasts for a minimum of three years, it typically takes five years, and in some cases it may require a minimum of seven

years.²⁰ Note that earnings rank and cost of living after the end of the residency would be part of the error term.

B. Instrumental Variables Model

The second model exploits the variation in beliefs induced by the source-randomization experiment to estimate the causal effects of perceptions on choice. Let $ER_{1,2}^{i,shown}$ be the information randomly chosen to be shown to the individual and $ER_{1,2}^{i,alt}$ be the alternative information that could have been shown to the individual but was not. Let $\Delta ER_{1,2}^i = ER_{1,2}^{i,shown} - ER_{1,2}^{i,alt}$ be the difference between the information shown and the alternative information that could have been shown. We estimate an IV-probit model that uses $\Delta ER_{1,2}^i$ and $\Delta COL_{1,2}^i$ as instrumental variables. In other words, this model uses the variation introduced by the random assignment of sources to estimate the effect of perceptions on choice:

$$\begin{aligned} I(Program_1 \succ_i Program_2) &= I(\beta^{ER} \times ER_{1,2}^{i,posterior} \\ &+ \beta^{COL} \times COL_{1,2}^{i,posterior} + \lambda_1 \times ER_{1,2}^{i,alt} + \lambda_2 \\ &\times COL_{1,2}^{i,alt} + \theta X^i + \varepsilon_i \geq 0), \\ ER_{1,2}^{i,posterior} &= \gamma_1^{ER} \times \Delta ER_{1,2}^i + \gamma_2^{ER} \times \Delta COL_{1,2}^i + \gamma_3^{ER} \\ &\times ER_{1,2}^{i,alt} + \gamma_4^{ER} \times COL_{2,1}^{i,alt} + \gamma_5^{ER} X^i + \varepsilon_{1,i}, \\ COL_{1,2}^{i,posterior} &= \gamma_1^{COL} \times \Delta ER_{1,2}^i + \gamma_2^{COL} \times \Delta COL_{1,2}^i \\ &+ \gamma_3^{COL} \times ER_{1,2}^{i,alt} + \gamma_4^{COL} \times COL_{2,1}^{i,alt} + \gamma_5^{COL} X^i + \varepsilon_{2,i}. \end{aligned}$$

There is a simple way to understand the intuition behind this instrumental variables approach. In a deceptive design, subjects would be shown the statistic from a certain source but with random noise added to this statistic. Then we would only exploit the variation in beliefs generated by the random noise. In our context, $\Delta ER_{1,2}^i$ and $\Delta COL_{1,2}^i$ play the role of the random noise added to the feedback, only they are generated in a nondeceptive manner.

IV. Implementation Details and Summary Statistics

In December 2016, we contacted the associate dean of student affairs at all 135 accredited medical schools in the United States by email to ask for permission to invite fourth-year students participating in the 2017 Main Residency Match to take part in our study (a sample of the invitation email is shown in appendix C.3). Our goal was to recruit as many respondents as possible, so we followed up, by email and phone, with all the deans who showed interest. Of the 79 schools that answered our invitation, 27 agreed to participate. The main

¹⁹The source for the demographic characteristics is the 2011–2014 American Community Survey. For the share of residents who are Democrats, we use the share of Obama voters between all voters in the 2008 presidential elections.

²⁰A small minority of subjects may expect to continue living in the same city after the residency, in which case, the distribution of earnings and the cost of living may also be relevant for the postresidency period.

reason given by the schools that declined to participate was school policy restricting external surveys, in place to avoid survey fatigue. Our sample of participating schools includes 22 of the 50 U.S. states, and it is quite representative of the whole sample of 135 accredited medical schools. We do not find statistically significant differences in observable characteristics such as total enrollment, average MCAT scores, undergraduate GPA at admission, acceptance rate, and *U.S. News and World Report* rank.²¹

For confidentiality reasons, we were not given email lists to directly invite students to participate in our study. Instead, the deans agreed to forward our invitation email containing the link to the survey to eligible students (i.e., senior medical students participating in the NRMP). This email invitation, a sample of which is shown in appendix C.4, asked students to participate in a confidential survey about the Main Residency Match for a study on how medical students select residency programs. The message mentioned that the survey would take less than ten minutes to complete and respondents would be sent a \$10 Amazon gift card by email as a token of appreciation. Finally, the email stressed the eligibility criteria for participating in the survey: being a graduating medical student participating in the Main Residency Match who has not yet submitted his or her rank to the NRMP.²²

The only reason we excluded individuals who had previously submitted their ranks was that we wanted individuals who were still deciding and thus prone to using the signals from the information-provision experiment. However, this concern is not important in the sense that submissions can be modified anytime before February 22. Even if some students had already submitted their rank at the time of responding to the survey, they would still be able to modify it. In any case, the vast majority of our subjects responded to the baseline survey quite early in the submission period.

We took several measures to minimize the chance that noneligible students would participate in the survey. First, deans were asked to carefully forward the invitation to senior students participating in the Main Residency Match. This request was not an issue since such a mailing list already existed; targeted announcements were already being sent to this group during the semester regarding the match. Second, individuals were reminded of these restrictions in the invitation email and on the consent page of the survey. Third, the first questions of the surveys acted as filters; we asked what match the respondent was participating in and whether they had already submitted their ranks. If they responded with a match other than the Main Residency or “yes” to already submitting their rank, the survey ended there, and they were excluded from taking the survey again.²³

²¹For details, see appendix A.3.

²²There are a number of alternative matches for some specialties that have different deadlines from the Main Residency Match.

²³The survey platform blocks users from taking the survey again by using their IP address and cookies, although students could circumvent this restriction by opening the survey link from a different device.

At the end of the survey, respondents were required to submit their university email address to sign a statement claiming that they were medical students participating in the NRMP and they understood that we reserved the right to verify their status before making a payment. We were able to confirm the validity of 100% of respondents for a subset of schools. Given all the measures taken and the evidence obtained, we are confident that the survey data are of high quality.

The invitation emails were sent to students in a staggered way, with the first round of invitations sent on January 6, 2017, and the last round of invitations and reminders sent on February 7. We estimated that the student invitations were forwarded to around 3,676 students, with 1,080 finishing the baseline survey, implying an overall response rate of 29.38%. The median survey completion time was almost 9 minutes. At the end of the baseline survey, we included an attention check question that was passed by 96.4% of respondents. For transparency, we do not drop the group that did not pass the attention check; indeed, we do not drop any other group from the baseline sample.²⁴

On February 23, 2017, the day after the NRMP rank submission deadline, we sent respondents who participated in our baseline survey an invitation to participate in the survey. We offered participants an additional \$5 Amazon gift card for participating in this shorter follow-up survey. We closed the follow-up survey on March 12, one day before Match Week started (i.e., the time when the students find out where they are matched). The response rate to the follow-up survey was 90.62%. Moreover, the characteristics of the individuals who responded to the follow-up survey are similar to the characteristics of individuals who did not respond to the follow-up.²⁵ On average, students responded to the baseline survey 24.5 days (SD 12.9) before submitting their ranks and responded to the follow-up survey 13.9 days (SD 11.8) after submitting their ranks (for more details, see appendix A.3).

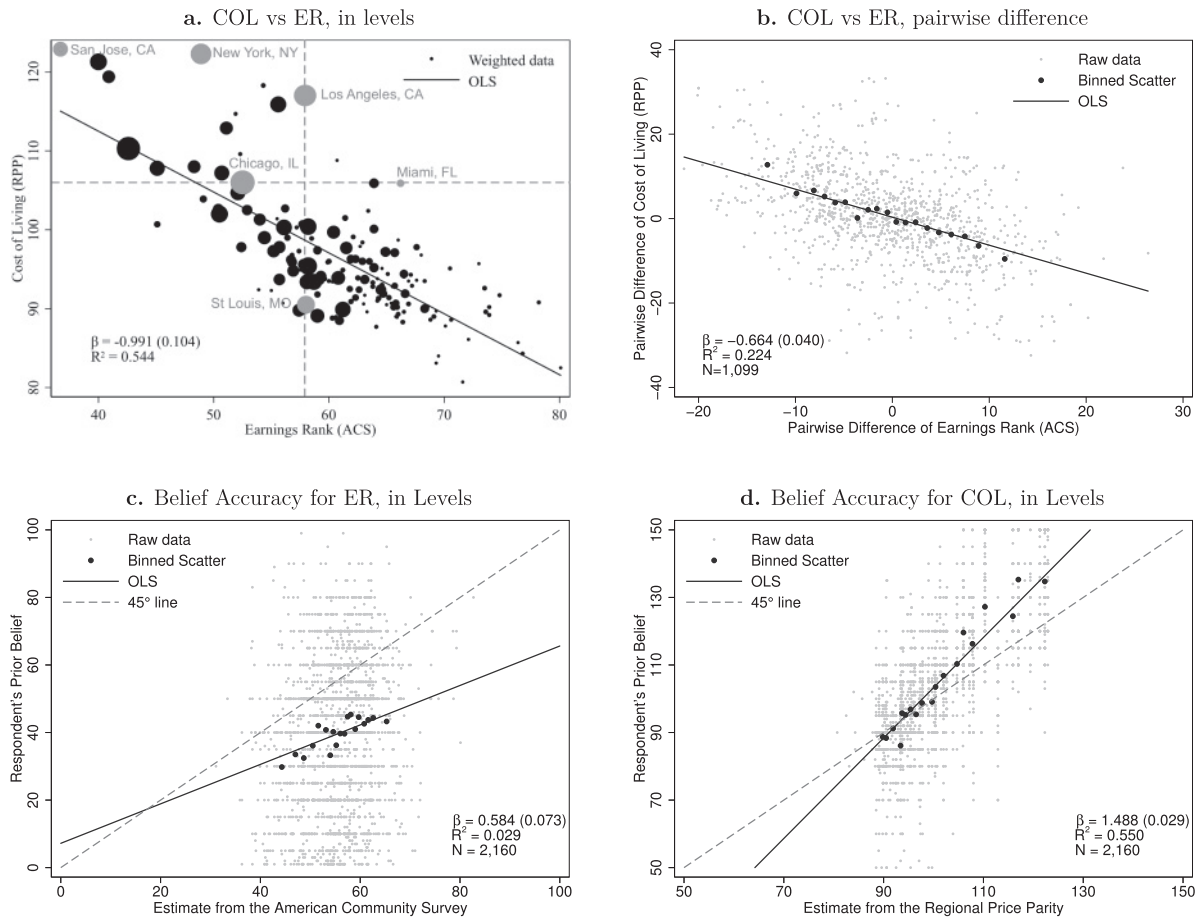
Appendix A.2 provides descriptive statistics about the subject pool. And consistent with successful random assignment, the characteristics are balanced across treatment groups. Around 48% of respondents were male, the average age was 27 years, 35.4% of respondents were single, 23.9% were married, and 40.7% were in a long-term relationship. On average, students were offered a salary of \$54,000 for the first year of their residency, a salary that would make them richer than 56% of earners in the average metro area. Of course, this sample is not representative of the general U.S. population of adults; most notably, our subject pool is younger and more educated. Nevertheless, our subject pool is close to the U.S. average in terms of nominal wages and gender composition.

The programs that the students are choosing over provide a broad geographical coverage of the U.S. territory (for more

²⁴Upon inspection of the data, the 3.6% of respondents who failed the attention check seem to have answered the survey as consistently as everyone else. And as reported in appendix A.9, the results are virtually the same if we drop this 3.6% of the sample.

²⁵Results presented in appendix A.3.

FIGURE 1.—VARIATION IN EARNINGS RANK, COST OF LIVING, AND BELIEF ACCURACY



(a) Dots are weighted responses for each metro area chosen by subjects. Variables for x-axis and y-axis correspond to levels of the earnings rank and the cost of living. (b) The gray dots correspond to the raw scatter plot, and the darker dots correspond to the binned scatter plot based on twenty bins. Variables for x-axis and y-axis correspond to pairwise differences across the two cities that the subject is considering submitting to the algorithm. Data from survey responses, the Regional Price Parity Index (for cost of living), and the American Community Survey (for earnings rank). (c, d) Comparison between respondent's perceptions before the information provision (i.e., prior beliefs) and statistics. The gray dots correspond to the raw scatter plot, and the darker dots correspond to the binned scatter plot based on twenty bins. Panels a and b present data in levels (i.e., two observations per individual, one for each of their options). Slopes (β , with robust standard errors in parentheses) and R^2 are based on a linear regression.

details, see appendix A.2). This broad coverage results in significant trade-offs between higher relative income and lower cost of living, where we observe significant orthogonal variation between them. Figure 1a shows the variation between the earnings rank (according to earning \$54,000 in the ACS) and cost of living (according to the RPP). There is substantial orthogonal variation in the two dimensions. For example, the metropolitan areas of Los Angeles and St. Louis, Missouri, have a similar earnings rank but are dramatically different in cost of living (difference of around 28 percentage points). On the other hand, figure 1b shows the orthogonal variation we use by plotting the pairwise differences in cost of living (according to the RPP measure) versus the differences in earnings rank (according to the ACS measure).²⁶ The substantial dispersion along the y-axis suggests that there are large differences in cost of living across the pairs of cities that the individuals must choose from. The substantial dispersion along the x-axis suggests large differences in earnings rank

across the pairs of cities.²⁷ Most important, the $R^2 = 0.22$ indicates that although the two are correlated,²⁸ substantial orthogonal variation exists between relative income and cost of living.²⁹

V. Results: Prior Beliefs and Learning

A. Distribution of Prior Beliefs

To the best of our knowledge, ours is the first paper to measure perceptions about earnings ranks and cost of living

²⁷Furthermore, the vast majority of these differences in earnings rank and cost of living are orthogonal to differences in nominal income. See appendix A.5 for details.

²⁸The slope of -0.664 suggests that, on average, relatively more expensive cities tend to have a higher distribution of nominal earnings.

²⁹We provide a number of additional checks in the appendix. Appendix A.10 shows that there is plenty of orthogonal variation between earnings rank and cost of living even when comparing residency programs of similar reputation. Appendix A.9 shows that the results are robust if we exclude individuals from the three largest metro areas (New York, Los Angeles, and Chicago) or if we exclude individuals from the 25 smallest metro areas (i.e., those with approximately fewer than 160,000).

²⁶Using the alternative data sources yields similar results.

across different cities. To get a sense of how informed individuals are about these aspects of their decision making, we start by comparing their prior beliefs (i.e., perceptions prior to the feedback) to the statistics from the baseline sources.

A growing literature documents that individuals have substantial misperceptions about their position in the national income distribution (Cruces et al., 2013; Karadja et al., 2017). Consistent with those findings, our data suggest that respondents have substantial misperceptions about their prospective earnings ranks. Figure 1c plots prior beliefs about earnings rank against our favorite estimates, from the ACS. On average, individuals underestimate earnings ranks by almost 16 percentage points. This finding is consistent with previous studies asking respondents to guess their position in the national income distribution. Karadja et al. (2017) find that Swedes underestimate their ranking by an average of 18 percentage points, which is close to the bias of 16 percentage points present in our own data. Moreover, the correlation between prior beliefs and our ACS estimates is positive and statistically significant (p -value < 0.001), but the R^2 is low (0.029). Expecting individuals to be accurate in predicting the levels of earnings ranks in different cities may be too demanding. Ultimately, the relevant margin for individuals choosing between two cities is the pairwise difference between the earnings ranks between those two places. Thus, we repeat this exercise using pairwise differences instead of levels of the earnings rank. Indeed, the misperceptions are less substantial when we look at these pairwise differences: the corresponding R^2 of the relationship between perceptions and the ACS estimates is 0.15.³⁰

In comparison, individuals are more accurate identifying the cost of living in the cities they are considering. Figure 1d shows respondents' prior beliefs about cost of living along with the corresponding RPP estimates. The RPP is meant to reflect all sources of expenditures, and for that, they employ data on prices in apparel, education, food, housing, medical, recreation, rents, transportation, and other goods and services. If prior beliefs were completely accurate, we would expect to see all responses on the 45 degree line. On average, prior beliefs overestimate the baseline estimate by just 4 percentage points, and the prior belief and RPP estimates are highly correlated, with an R^2 of 0.550.

B. Learning from Feedback

We next examine whether respondents learned from the information we provided. If they do learn, we would expect a positive relation between their perception gaps (i.e., the signal received minus the prior belief) and their revisions (i.e., the posterior belief minus the prior belief); that is, respondents who originally overestimated would revise their beliefs downward, while those who underestimated would revise in the opposite direction. Moreover, the slope of this relationship can be used to quantify the degree of learning

from the feedback: a parameter that takes the value from 0 to 1, indicating the weight assigned to the signal relative to the weight assigned to the prior belief.³¹

Respondents strongly updated their beliefs after being provided with feedback. Figure 2 presents the reduced-form effects of information for earnings rank and cost of living, respectively. Figures 2a and 2d present the short-term effect, that is, the revision made by respondents directly after being given the information. The short-term learning rates, given by the slopes reported in these figures, are 0.873 (SE 0.011) for the earnings rank and 0.879 (SE 0.010) for the cost of living. These two learning rates are statistically significant (both p -values < 0.001), precisely estimated, and we cannot reject the null hypothesis that they are equal to each other (p -value = 0.754). These learning rates are remarkably close to 1, meaning that respondents almost fully reacted to the signals.

One limitation with this evidence is that individuals may have revised their beliefs toward the truth regardless of the feedback we provided. For instance, they may have taken extra time to think about the question, leading to a more accurate response. The source experiment was designed to test this specific hypothesis. We construct two variables: the information actually shown and the "alternative" information that could have been shown. If the alternative information had any effect beyond the information shown, that would be evidence that some of the revisions were due to reversion to the truth rather than reversion to the information provided. Figures 2b and 2e show the relation between the alternative information and the revision adjusted for the information actually shown. The alternative information indeed has no effect: the coefficients are close to 0 (0.060 for earnings ranking and -0.034 for cost of living) and precisely estimated. Furthermore, in appendix A.19, we show that there was no cross-learning (i.e., feedback on cost of living did not affect beliefs about earnings ranking and vice versa).

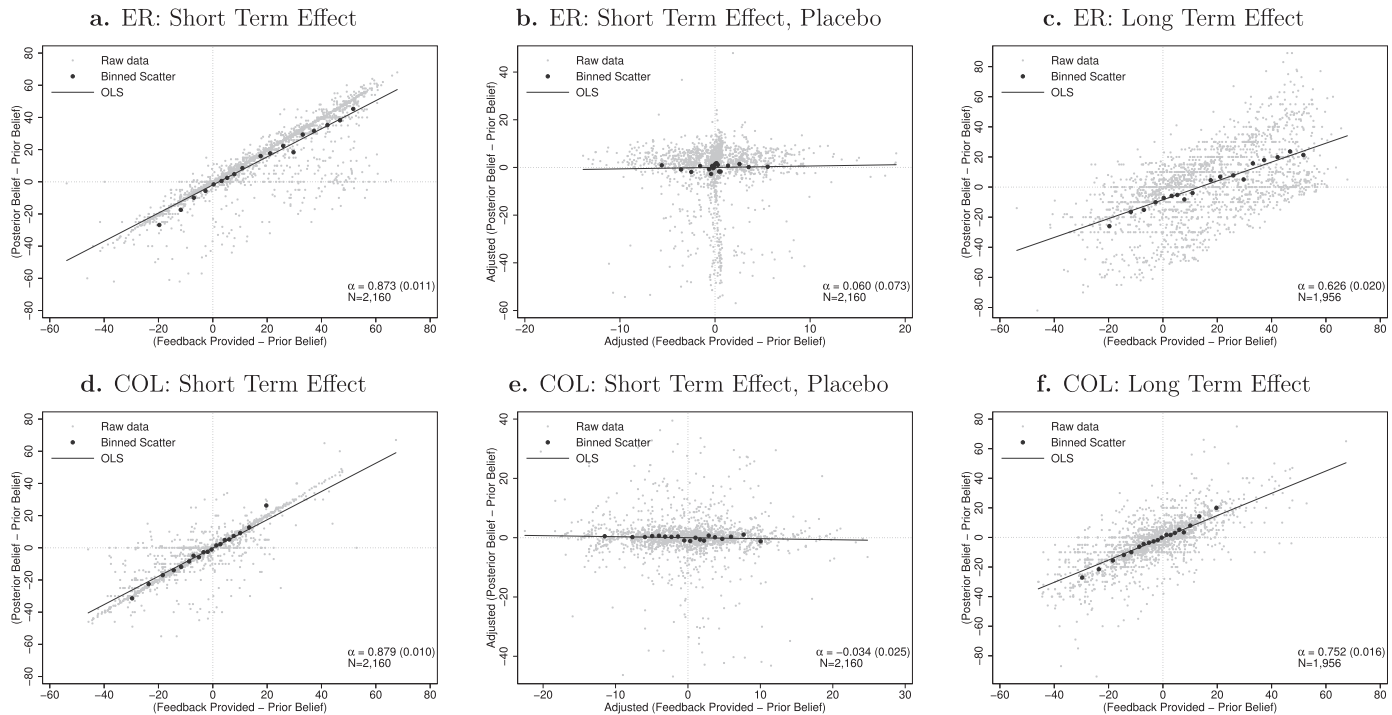
In survey experiments, one main concern is that instead of inducing genuine learning, the information provided in the experiment may elicit spurious reactions. For instance, if an individual is told that they would be "richer than 60%" of that city's population and then later asked about the cost of living in the same city, he or she may report an earnings rank that is closer to 60% for spurious reasons, such as unconscious numerical anchoring (Kahneman & Tversky, 1972). Under the assumption that these effects are temporary, we can disentangle genuine from spurious learning by looking at the reaction to the information provided in the experiment that persisted over time (Cavallo, Cruces, & Perez-Truglia, 2017).

We look at the persistence of the effect of feedback between the time participants responded to the baseline and follow-up surveys, which was 38.4 days on average. Figures 2c and 2f show the relation between the initial perception gap and the long-term revision based on beliefs reported in the follow-up survey. There is substantial persistence of the

³⁰Detailed results reported in appendix A.6.

³¹See appendix A.7 for more details.

FIGURE 2.—LEARNING FROM THE EXPERIMENTAL FEEDBACK



Comparison between the difference in statistics and respondent's perceptions before the information provision (i.e., prior beliefs), and difference in respondent's perceptions after the information provision (i.e., posterior beliefs) and prior beliefs. The gray dots correspond to the raw scatter plot, and the darker dots correspond to the binned scatter plot based on twenty bins. Panels b and e show a placebo test where we compare the difference between the alternative feedback and prior belief to the difference between the posterior and prior beliefs, adjusting for the shown statistic. Panels c and f use respondent's perceptions measured in the follow-up survey as posterior belief. The slope (α , with robust standard errors in parentheses) is based on a linear regression.

effects of the feedback: the estimated slope for the initial perception gap and the long-term revision (i.e., the difference between long-term belief and the initial prior belief) for earnings rank is 0.626 (SE 0.020), while for cost of living it is 0.752 (SE 0.016). The persistent effect of our feedback suggests that individuals indeed care about these city features. These longer-term revisions are slightly weaker than the short-term revisions, but that result is expected given that individuals must have gathered some additional information in the time between the two surveys.

VI. Results: Location Preferences

A. Average Preferences

We start with the baseline specification, which uses the probit model from section III with the expected rank submission as the dependent variable. This specification exploits all the variation in perceptions, which includes the experimental variation induced by our information provision as well as the remaining nonexperimental variation. We introduce variations of this specification later.

The probit coefficients are presented in table 1. Column 1 presents the results for the full sample, while columns 2 through 7 present results by demographic subgroups. The average subject also prefers a higher earnings rank: the estimated β^{ER} from column 1 is positive and statistically significant (p -value = 0.065). This coefficient suggests that the aver-

age individual prefers to live in a city where, holding her cost of living constant, she earns more than her neighbors. To better understand the magnitude of these probit coefficients, we can transform them into the corresponding marginal effects, where increasing the earnings rank at a program's location by 1 percentage point increases the probability of choosing that program by 0.186 percentage points (for a behavioral elasticity of 0.186).³²

We use preferences over cost of living as a benchmark for preferences for relative income. The estimated β^{COL} from column 1 is negative and statistically significant (p -value = 0.027), suggests that the average individual prefers programs with lower costs of living. The corresponding marginal effect indicates that increasing the cost of living by 1 percentage point at a program's location decreases the probability of choosing that program by 0.201 percentage points (which can be interpreted as a behavioral elasticity of -0.201). The elasticity for cost of living (-0.201) is similar in absolute value to the elasticity for earnings rank (0.186); indeed, their difference is statistically insignificant. This finding suggests that individuals care about relative income nearly about as much as they care about cost of living: the average individual would be willing to accept an increase of 0.925 percentage points in the cost of living in exchange for a 1 percentage point increase in earnings rank.

³²These marginal effects are reported in appendix A.8.

TABLE 1.—LOCATION PREFERENCES: BASELINE ESTIMATES

	By Relationship Status		By Gender		By Specialty Salary		
	All (1)	Nonsingle (2)	Single (3)	Female (4)	Male (5)	≤ \$229,000 (6)	> \$229,000 (7)
β^{ER}	0.995* (0.539)	2.236*** (0.669)	-1.538* (0.880)	1.041 (0.755)	0.896 (0.781)	1.433* (0.732)	0.777 (0.797)
β^{COL}	-1.073** (0.485)	-1.087 (0.663)	-1.058 (0.749)	-0.972 (0.679)	-1.443* (0.753)	-0.690 (0.713)	-1.238* (0.690)
Diff. <i>p</i> -value [<i>q</i> -value]							
ER		0.001 [0.030]		0.894 [0.974]		0.544 [0.954]	
COL		0.977 [0.977]		0.642 [0.954]		0.580 [0.954]	
Observations	1,080	698	382	560	520	549	531

Heteroskedasticity-robust standard errors in parentheses. *** *p* < 0.01, ** *p* < 0.05, and * *p* < 0.1. Raw probit coefficients. Each column corresponds to a different probit regression of expected rank order submission on posterior beliefs about earnings rank and cost of living, from the baseline survey, including the baseline controls listed in section III. Columns 2 through 7 show estimates when restricting sample to different subgroups: columns 2 and 3 by nonsingle (i.e., married or in a long-term relationship) or single, columns 4 and 5 by gender, columns 6 and 7 by expected salary of specialty after residency (over and below the median value of \$229,000). *P*-values correspond to the test of the null hypothesis that the coefficients are equal between the two subgroups, multiple-testing *q*-values based on Benjamini and Yekutieli (2001) presented in brackets.

The fact that medical students care about relative income and cost of living is consistent with the view that money is a primary motivation for doctors. For instance, according to a 2008 survey, 49% of pre-med students self-reported being primarily motivated by money in their career choice (Daniel & O’Brien, 2008). Our findings are also consistent with prior survey evidence indicating that NRMP applicants care about the geographic location of the programs. For example, Sledge et al. (1989) surveyed a sample of graduating medical students to elicit the most important features in determining the program they ranked first in their NRMP submission. They allowed subjects to choose from fifteen features that included characteristics of the program (e.g., program curriculum) and other characteristics (e.g., geographic location). They found that geographic location was rated the most important feature (selected 63.8% of the time) as well as the most desirable feature (selected 37.6% of the time). Using more recent data, NRMP (2015) also find that geographic location is the most popular feature; they also find that when asked to mention the most important factors taken into account in their decision making, 52% of surveyed medical students explicitly mention the cost of living.

Although β^{ER} and β^{COL} are statistically and economically significant, they do not imply that relative income and cost of living during the residency are the main features that medical students pay attention to. Indeed, if doctors are motivated by financial rewards, we would expect them to seek a higher postresidency income through more prestigious programs. Indeed, in complementary analysis, we find that when deciding their NRMP submissions, doctors care more about the prestige of the program and the career prospects than about the relative income and cost of living during their residency. For example, a standard deviation increase in sense of purpose of a program has an effect that is 3.9 times larger than a standard deviation increase in earnings rank (results presented in appendix A.11). This finding is also consistent with survey data indicating that the program’s prestige is of higher importance than other attributes such as cost of living (NRMP, 2015).

Under the assumption that our revealed-preference estimates reflect the same type of neighbor externalities from Luttmer (2005), we can compare the magnitude of relative concerns in this paper to the magnitude reported in Luttmer (2005). This comparison, presented in appendix A.13, suggests that our estimates imply a somewhat smaller role for relative concerns.³³

B. Heterogeneity by Relationship Status

The average preferences could potentially mask substantial heterogeneity. For instance, Luttmer (2005) finds that the effect of relative income on happiness is driven entirely by non-single individuals. Also, evidence from the urban economics literature indicates that single and nonsingle individuals have different location preferences (Couture & Handbury, 2016; Gautier, Svarer, & Teulings, 2010). With the goal of exploring this form of heterogeneity, we elicited the relationship status using the same categories as in Luttmer (2005).

Columns 2 and 3 of table 1 present the results on this form of heterogeneity by estimating the model separately for singles and nonsingles. Column 2 of table 1 shows the estimates for nonsingle individuals (i.e., the 65% of the sample who are married or in a long-term relationship) and column 3 for the sample of single individuals (35% of the sample).³⁴ It is important to note that by “nonsingle,” we refer only to their relationship status, not to whether the respondent participates as a dual match, which is a special regime used by roughly 7% of subjects; indeed, the results are similar if

³³For a more direct comparison between happiness and choice data, our survey collected a question on happiness: “If assigned to it, in which of the two programs would you expect to live a happier life?” Responses used the same likelihood scale as for the rank choices. We find that the marginal rates of substitution inferred by happiness are statistically indistinguishable from the marginal rates of substitution inferred by choice; however, due to the lack of precision of the happiness estimates, we cannot rule out large discrepancies. Results are presented in appendix A.14.

³⁴Appendix A.12 shows results breaking down the nonsingle individuals into married and in a long-term relationship. The relative concerns are similar between these two groups.

we drop the small minority of subjects participating in dual matches.³⁵

The coefficients from columns 2 and 3 suggest a large heterogeneity in β^{ER} by relationship status. For nonsingle individuals, the estimated β^{ER} (2.236) is positive and highly statistically significant (p -value = 0.001). For the sample of single individuals, β^{ER} (−1.538) is negative and statistically significant at the 10% level (p -value = 0.080). The direction of the difference in relative concerns between nonsingles and singles is consistent with the evidence from Luttmer (2005). The difference in β^{ER} between nonsingles and singles is also highly statistically significant (p -value = 0.001). And contrary to the case of preferences for earnings rank, the relationship status does not seem to affect the preferences for cost of living. According to columns 2 and 3 of table 1, our benchmark, β^{COL} , is −1.087 for nonsingles and −1.058 for singles, with the difference being statistically insignificant (p -value = 0.977).

These estimates suggest that while nonsingle individuals prefer to live in less affluent ponds, single individuals would rather live in more affluent ponds. In other words, single individuals tend to see more affluent neighbors as a positive rather than a negative externality. One potential explanation is that they expect to benefit from some public goods that are provided by the more affluent neighbors. Another potential explanation lies in the dating market: since more affluent individuals prefer to date more affluent partners (Fisman et al., 2006; Hitsch, Hortaçsu, & Ariely, 2010), this can naturally create a preference for locating in more affluent ponds.³⁶ Indeed, this result is consistent with Gautier et al. (2010), who show that singles are willing to pay higher housing prices to benefit from a denser dating market in cities. Consistent with our findings for nonsingles, Gautier et al. (2010) find that after getting married, the dating-market benefits no longer matter for them, and couples move out of the city.

It is important to note that the singles in our subject pool are a special group of individuals, and as a result, the dating motive in this sample may be stronger than for the general population of singles. First, these subjects are at their prime dating age, and thus they may expect to find a long-term partner during their residency. Second, although their wages during the residency put them near the middle of the national earnings distribution, their expected *postresidency* earnings will place them at the very top of the national earnings distribution. As a result, these subjects would probably struggle to meet a partner who can match their permanent income unless they move into one of the most affluent cities.

Columns 4 through 7 of table 1 explore other forms of heterogeneity based on characteristics measured in the base-

line survey. Columns 4 and 5 explore potential differences in preferences by gender: β^{ER} is similar for females (1.041) and males (0.896), and β^{COL} is also similar for females (−0.972) and males (−1.443). Moreover, neither of these two differences is statistically significant (p -values of 0.894 and 0.642, respectively). Columns 6 and 7 explore heterogeneity by future income. Although all of these subjects receive a similar income during the residency, they have different expected incomes after they finish their residencies. Then, it is possible that individuals who selected high-earning specialties may be more concerned about cost of living or relative income. To test this hypothesis, columns 6 and 7 split the sample in specialties with above- and below-median postresidency average income. The β^{ER} is 1.433 for below-median specialties and 0.777 for above-median specialties, and β^{COL} is −0.690 for below-median specialties and −1.238 for above-median specialties, with neither of those differences being statistically significant (p -values of 0.544 and 0.580, respectively).

To address the possibility of false positives due to multiple hypothesis testing, for each p -value reported in table 1, we report the corresponding q -value in brackets (Benjamini & Yekutieli, 2001). The q -value indicates the minimum false discovery rate (i.e., the expected proportion of rejected null hypotheses that are actually true) at which the null hypothesis would be rejected for that test given all tests reported in the same table. The adjustment for multiple hypothesis testing does not change any of the results: most important, the difference in β^{ER} between singles and nonsingles has a q -value of 0.030, which indicates that this heterogeneity is unlikely to be spurious.

We also computed heterogeneity by other characteristics measured in the follow-up survey. None of these dimensions are nearly as important as relationship status for predicting heterogeneity in preferences for relative income, in terms of both magnitude and statistical significance.³⁷ Because of the magnitude of the heterogeneity by relationship status, in the remainder of the paper, we report estimates for the entire sample, as well as broken down by relationship status.

C. Robustness Checks: Additional Controls

One potential concern with the baseline specification is that of omitted-variable biases. For instance, if places where an individual expects higher earnings rank (i.e., less affluent metro areas) are systematically worse in terms of quality of life, then failing to account for quality of life would introduce a negative bias in β^{ER} . This would make relative concerns look weaker than they actually are.

Table 2 presents the robustness checks. Each row corresponds to a different regression, with a different set of control variables. The first row presents results for our baseline specification but without including any control variables for the characteristics of the program or the metro area. The second row corresponds to the baseline specification from table 1,

³⁵See appendix A.9 for more details.

³⁶Prior evidence suggests that relative to single men, single women may have a stronger preference for finding more affluent partners (Bertrand, Kamenica, & Pan, 2015; Bursztyn et al., 2017). Consistent with this view, we find that the preference for more affluent ponds among singles is driven primarily by single women, although this result is imprecisely estimated. Results reported in appendix A.12.

³⁷Results are reported in appendix A.18.

TABLE 2.—LOCATION PREFERENCES: ROBUSTNESS TO ALTERNATIVE CONTROL VARIABLES

	A: β^{ER}			B: β^{COL}		
	Nonsingle (1)	Single (2)	All (3)	Nonsingle (4)	Single (5)	All (6)
No Controls	1.961*** (0.663)	-1.480* (0.841)	0.873* (0.531)	-0.812 (0.523)	-1.131* (0.589)	-0.894** (0.382)
Baseline	2.236*** (0.669)	-1.538* (0.880)	0.995* (0.539)	-1.087 (0.663)	-1.058 (0.749)	-1.073** (0.485)
Demographic	2.288*** (0.715)	-0.871 (0.977)	1.176** (0.578)	-1.219* (0.628)	-1.712** (0.713)	-1.342*** (0.468)
Amenities	2.056*** (0.669)	-1.381 (0.853)	0.958* (0.538)	-0.718 (0.630)	-1.265 (0.816)	-0.898* (0.481)
Geography	2.011*** (0.658)	-1.214 (0.869)	1.010* (0.528)	-1.093** (0.547)	-1.393** (0.592)	-1.168*** (0.401)
Economic	1.914*** (0.684)	-1.191 (0.941)	0.946* (0.566)	-0.467 (0.670)	-1.647** (0.812)	-0.868* (0.498)
State Dummies	2.901*** (0.703)	-1.943** (0.907)	1.084* (0.555)	-1.090 (0.671)	-1.219 (0.939)	-0.968* (0.502)
Obj. Program Chars.	2.018*** (0.676)	-1.578* (0.847)	0.904* (0.535)	-0.840 (0.516)	-1.151** (0.587)	-0.910** (0.381)
Subj. Program Chars.	2.222*** (0.730)	-1.320 (1.040)	1.199** (0.605)	-1.210** (0.587)	-1.678*** (0.619)	-1.277*** (0.425)
All Controls	2.385*** (0.887)	-0.685 (1.797)	1.147* (0.684)	-0.550 (1.008)	-4.165*** (1.598)	-1.097 (0.686)

Heteroskedasticity-robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$. Raw probit coefficients. Each row corresponds to a separate regression of expected rank on posterior beliefs about earnings rank and cost of living, from the baseline survey. All regressions include as controls the log difference in nominal income and a constant. The first row does not include any additional controls. The second row includes the baseline controls listed in section III. The third-to-last rows use different sets of additional controls, listed in section VI.C. Results are based on 1,080 individual responses (698 from nonsingles and 382 from singles), except for the last two rows, which are restricted to the follow-up sample (978 responses; 595 from nonsingles; 311 from singles).

which includes the six baseline controls listed in section III.A. The results in the first two rows of table 2 indicate that β^{ER} and β^{COL} are qualitatively and quantitatively similar between the baseline specification and the specification without controls. The third through last rows of table 2 offer different sets of additional controls. These sets of controls were selected based on attributes that could potentially be relevant and at the same time may be correlated to the earnings rank. For instance, if the best rated residency programs were located in less affluent cities, that could bias our estimated preferences for relative income.

Each row introduces a different set of control variables: demographic characteristics (population, population density, percentage urban population, percentage same gender, percentage age 25 to 34, share of college graduates, share foreign, share Hispanic, and share black); amenities (quality of life from Albouy, 2016, per capita spending on local public goods, per capita spending on education and health, overall crime rate and violent crime rate, share of registered Democrat voters in the 2012 election); geography (distance of program to current medical school); economic factors (estimated income taxes, federal and state income taxes, local sales tax, rent prices, and the Gini coefficient); a set of state dummies; objective program characteristics (residency program rank from Doximity, dummies for university hospitals); and subjective program characteristics (the subjective rank in prestige, purpose, and prospect, as reported in the follow-up survey).³⁸ Comparing the results across rows of β^{ER} and β^{COL} of

³⁸The results are similar when using additional controls related to the availability and quality of residency programs in the area. Results are reported in appendix A.10.

table 2 suggests that these estimates are robust to the choice of control variables in terms of both statistical significance and economic significance.³⁹

D. Robustness Checks: Experimental Estimates and Persistence

In this section, we present results from two robustness checks. The first check addresses concerns about omitted-variable bias by exploiting the exogenous variation in beliefs generated by the source-randomization experiment. The second is intended to address potential concerns about spurious effects of the information-provision experiment by comparing the short-term effects to the long-term effects of the information provision. To make these estimates directly comparable to the long-term effects, in this section we restrict the sample to the 90.6% of subjects who responded to the follow-up survey.

Panel A of table 3 presents the results for β^{ER} . The first row presents the baseline specification, while the second row presents the experimental estimates. The experimental estimates are less precisely estimated than the baseline estimates

³⁹Of course, small differences occur in the point estimates across specifications. For instance, relative to the baseline β^{ER} of 0.995 for the entire sample in column 3, β^{ER} ranges from a minimum of 0.873 with no controls to a maximum of 1.199 with subjective program characteristics. However, all of these differences are statistically insignificant. Also, according to the pseudo- R^2 , including these variables increases the explanatory power of our model to some degree. For the full sample, the pseudo- R^2 increases from 0.015 in the specification with no additional controls to a minimum of 0.018 with controls for objective program characteristics or amenities, and a maximum of 0.123 with controls for subjective programs characteristics (or 0.218 with all controls).

TABLE 3.—LOCATION PREFERENCES: EXPERIMENTAL ESTIMATES

	A: β^{ER}			B: β^{COL}		
	Nonsingle (1)	Single (2)	All (3)	Nonsingle (4)	Single (5)	All (6)
Baseline	2.380*** (0.702)	-1.656* (0.991)	1.141** (0.577)	-1.234* (0.743)	-1.379* (0.772)	-1.262** (0.531)
Experimental	2.977** (1.331)	-4.964** (1.974)	0.867 (1.151)	0.353 (1.160)	1.663 (1.286)	0.662 (0.881)
Experimental, Long Term	1.993* (1.188)	-5.285*** (1.984)	-0.029 (1.071)	1.662* (1.005)	0.251 (1.359)	1.012 (0.821)
Experimental, Falsification	-0.007 (0.998)	0.040 (1.732)	0.004 (0.837)	0.037 (0.855)	0.021 (1.123)	0.031 (0.651)

Heteroskedasticity-robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$. Raw probit (or IV-probit) coefficients restricting sample to respondents who completed the follow-up survey. All regressions include the baseline controls listed in section III. The independent variables are the posterior beliefs about earnings rank and cost of living, from the baseline specification. The first row corresponds to the baseline probit specification. The second through third rows correspond to IV-probit regressions, using the variation in perceptions generated by the source-randomization experiment as instrumental variables. The first and second rows use the expected rank-order submission (from the baseline survey) as the dependent variable. The third row uses the final rank-order submission (from the follow-up survey) as the dependent variable. The fourth row corresponds to a falsification test that uses the same IV-probit specification from the second row, but using the list order (i.e., the order in which programs are listed at the beginning of the survey) as the dependent variable instead of the rank order. To estimate this IV-probit model, we randomly assign programs to be program 1 and program 2, and then we use as dependent variable a dummy that takes the value 1 if program 1 was listed first at the beginning of the survey. We repeat this procedure 1,000 times and report the average and standard error from the distribution of coefficients. All results based on the sample of individuals who completed the follow-up survey (978 responses; 647 from nonsingles; 311 from singles).

because they use only a portion of the variation in beliefs. For each of the subgroups of single and nonsingle respondents, shown in columns 1 and 2, the estimated β^{ER} is qualitatively consistent across the baseline and experimental specifications. For nonsingles, the coefficient is 2.380 (p -value = 0.001) in the baseline specification versus 2.977 (p -value = 0.025) in the experimental specification. And for singles, the coefficients are -1.656 (p -value = 0.095) in the baseline specification versus -4.964 (p -value = 0.012) in the experimental specification.

Column 3 shows that for the entire sample, β^{ER} is slightly lower in the second row (0.867) than in the first row (1.141) and, due to the lower precision, becomes statistically insignificant in the second row. However, we must take this finding with a grain of salt. First, due to the precision of the experimental coefficient, this difference between the first and second rows is statistically insignificant. Second, the reduction in the average β^{ER} is driven primarily by the fact that the coefficient becomes more negative (and thus more consequential) among singles.

Panel B of table 3 presents the results for β^{COL} . The results from the baseline specification (first row) are qualitatively different from the results in the experimental specification (second row). All of the coefficients (for the entire sample, singles and nonsingles) become positive, are imprecisely estimated, and are statistically insignificant. One potential interpretation could be that individuals anticipate that unobservable city amenities translate into higher cost of living, thus reducing the desire to move to areas with lower cost of living (Roback, 1982; Albouy, 2008). However, this may simply be a spurious finding: since the experimental estimates are not precisely estimated, we cannot rule out large, negative values for β^{COL} , and in most cases, we cannot reject that the experimental coefficients are equal to those from the baseline specification.⁴⁰ Indeed, consistent with the interpre-

tation of spurious finding, we find a negative and significant β^{COL} when we replicate the experiment with the secondary sample (results reported in section VI.E).

As discussed above, the treatment groups were balanced in observable characteristics, suggesting that the randomization was indeed successful. As an additional robustness check, we reestimate the instrumental variables model, but instead of the rank order, we use the list order as the dependent variable (i.e., the order in which the individual listed the residency programs at the beginning of the survey). Because it takes place before the provision of feedback, the feedback should not have any effect on the list order. We present results for this falsification in the fourth row of table 3. As expected, the estimated values of β^{ER} and β^{COL} are close to 0 and statistically insignificant, in the full sample as well as in the subsamples of nonsingles and singles.

The appendix presents some additional results. In all of the instrumental variable specifications, we strongly reject the null hypothesis of weak instruments. Also, the learning rates implied by the first-stage coefficients are always close to 1, and for that reason, the instrumental variables estimates are similar to the reduced-form estimates.⁴¹

The second robustness test is intended to address potential concerns about spurious effects of the information-provision experiment. By asking individuals questions about the earnings rank and cost of living, the baseline survey makes those aspects more salient, which may make individuals overweight them in their expected choice. However, it must be noted that this salience effect may not necessarily exaggerate the importance of relative income because it would be expected to inflate both β^{ER} and β^{COL} . Another potential concern is that of experimenter-demand effects: by providing information about earnings rank and cost of living, the experimenter

average preferences instead of average preferences. For instance, it is possible that the information-provision experiment disproportionately affected individuals who were the most unsure about their prior beliefs about cost of living, who likely are those who care the least about cost of living.

⁴¹Reduced-form and first-stage estimates are presented in appendix A.18.

⁴⁰Also, the coefficients from the first and second rows are not-expected to be equal to the extent that the experimental coefficients identify local

may be putting pressure on the subjects to use this information in their expected choice. Again, this source of bias would not necessarily exaggerate the importance of relative income: since most individuals do not want to reveal to others that they care about status (Shigeoka & Yamada, 2016), the experimenter-demand bias would probably shrink β^{ER} toward 0.⁴²

To address these remaining concerns, we estimate the effects of the information provision on the final rank submission, which takes place an average of 38.4 days after the information provision. This can be achieved by using the same instrumental variable model, but using the final submission rank (elicited in the follow-up survey) instead of the expected submission rank (elicited in the baseline survey) as the dependent variable. If the effects were purely due to salience or experimenter-demand effect, we would expect that the information provided in the experiment would not have any effect on the final submission choice.

The third row of table 3 presents the experimental estimates based on the long-term effects of the experiment. By comparing the coefficients in the third row to those from the second row, we can compare the short-term and long-term effects of the information. The long-term experimental coefficients are somewhat different from the short-term experimental coefficients, but those differences are mostly statistically insignificant. Most important, the coefficient on β^{ER} is still positive (1.993) and statistically significant at the 10% level (p -value = 0.093) for nonsingles, and negative (−5.285) and statistically significant at the 1% level (p -value = 0.008) for singles.

E. Robustness Checks: Replication

In appendix B, we provide an additional robustness check, based on an auxiliary experiment conducted with a sample of respondents recruited through Amazon Mechanical Turk. This sample has the disadvantage that respondents are not planning to move anytime soon, so we must study their hypothetical location choices instead of their actual choices. The auxiliary sample does have some advantages, however: it is more diverse than the primary sample of medical students in many dimensions such as age and occupation, and it is possible to run further experiments at any time. The results from this auxiliary experiment replicate the main findings from the original experiment: we find that the average individual prefers a higher relative income and a lower cost of living; we find that the experimental estimates are similar to the nonexperimental estimates; and we find large differences in preference for relative income by relationship status. We also compare the marginal rate of substitution between relative income and cost of living between the auxiliary and main

samples. The results indicate that while somewhat smaller in magnitude, the preference for relative income is still statistically and economically significant in the, more diverse, auxiliary sample.

F. Robustness Checks: Alternative Mechanisms

One possible interpretation of the coefficient β^{ER} is that individuals use their prospective relative income as a signal for other city attributes. Given that medical students are highly informed and are making a high-stakes decision, this possibility seems unlikely. These students devote their entire fourth year of medical school to the match. After the hospital visits, they have about two months to finalize their rankings. During this time, they continue to gather information to aid their decision. Because they have been to these locations and can easily obtain additional information directly, it is unlikely that they would rely on earnings rank to learn about other features of the locations. In any case, we present two tests for this mechanism. One version of this mechanism is that individuals use information about relative income to make inferences about the expected cost of living. If participants in the NRMP believe that their earnings rank reflects the degree of competition with their neighbors for some goods, such as housing, it would be natural for them to learn about cost of living from information about their relative income. In appendix A.19 we provide evidence against this confounding factor by showing that the information about relative income did not affect posterior beliefs about the expected cost of living. A second version of this mechanism is that individuals use information about relative income to learn about other city characteristics, such as school quality and crime rates. In appendix B.6, we use data from the auxiliary experiment, in which we elicited posterior beliefs about these and other attributes of the city, to show evidence against this confounding factor.

VII. Conclusion

We conducted a field experiment with a sample of 1,080 medical students who participated in the National Resident Matching Program. We provide unique revealed-preference evidence that individuals care about their prospective relative income. We find that, on average, individuals want to live in places where their relative income will be higher. This evidence is consistent with the predictions of models of relative concerns, according to which affluent neighbors would impose a negative externality (Luttmer, 2005). We also show that there is meaningful heterogeneity in these preferences for relative income: nonsingle individuals want to live in less affluent ponds, whereas single individuals prefer to live in more affluent ponds. The heterogeneity by relationship status suggests that relative to nonsingles, singles see more affluent neighbors as a positive externality. For example, singles may expect to benefit from dating their affluent neighbors or from the public goods that they finance.

⁴²Also, our survey was conducted confidentially and online, which reduces the scope for experimenter-demand effects (Van Gelder, Bretveld, & Roeleveld, 2010). Additionally, it would be difficult to reconcile the experimenter-demand channel with the finding that the earnings rank had a positive effect on nonsingles and a negative effect on singles.

Regarding the external validity of our results, it is possible that senior medical students care about their relative income more than the general population does. For example, given that most doctors make well above the subsistence level, they may have the luxury of caring about their relative income. Doctors may also be more materialistic (Daniel & O'Brien, 2008). Indeed, the results from our auxiliary survey, based on a more diverse cross-section of the U.S. population, suggests that the preference for relative income is still present but somewhat smaller in magnitude for the general population. An avenue for future research is to find other contexts in which this revealed-preference method could be used to estimate preferences for relative income. For instance, although the settings may not be as clear-cut as for the medical residency, there are multiple job markets in which job seekers must choose between job offers in different cities.⁴³ Future research could also investigate the mechanisms underlying individuals' concerns about relative income. There is, for instance, little evidence as to whether the preference for relative income responds to instrumental (e.g., dating prospects) versus noninstrumental motives (e.g., envy).⁴⁴ We hope our methodology could be used in the future to study not only preferences for relative income and cost of living but also preferences over various different city amenities.

⁴³Some examples are employees on the job market considering job offers, prospective undergraduate or graduate students choosing where to study, and graduating students choosing among their jobs after graduation.

⁴⁴These additional hypotheses can be explored with the same empirical framework proposed in this paper, but with additional treatment arms designed to test specific mechanisms.

REFERENCES

- Agarwal, N., "An Empirical Model of the Medical Match," *American Economic Review* 105:7 (2015), 1939–1978. 10.1257/aer.20131006, PubMed: 29543417
- Albouy, D., "Are Big Cities Bad Places to Live? Estimating Quality of Life across Metropolitan Areas," NBER working paper 14472 (2008).
- , "What Are Cities Worth? Land Rents, Local Productivity, and the Total Value of Amenities," this REVIEW 98:3 (2016), 477–487.
- Bayer, P., F. Ferreira, and R. McMillan, "A Unified Framework for Measuring Preferences for Schools and Neighborhoods," *Journal of Political Economy* 115:4 (2007), 588–638. 10.1086/522381
- Bénabou, R., and J. Tirole, "Incentives and Prosocial Behavior," *American Economic Review* 96:5 (2006), 1652–1678.
- Benjamin, D. J., O. Heffetz, M. S. Kimball, and A. Rees-Jones, "Can Marginal Rates of Substitution Be Inferred from Happiness Data? Evidence from Residency Choices," *American Economic Review* 104:11 (2014), 3498–3528. 10.1257/aer.104.11.3498, PubMed: 25404759
- Benjamini, Y., and D. Yekutieli, "The Control of the False Discovery Rate in Multiple Testing under Dependency," *Annals of Statistics* 29:4 (2001), 1165–1188. 10.1214/aos/1013699998
- Bertrand, M., E. Kamenica, and J. Pan, "Gender Identity and Relative Income within Households," *Quarterly Journal of Economics* 130:2 (2015), 571–614. 10.1093/qje/qjv001
- Black, S. E., "Do Better Schools Matter? Parental Valuation of Elementary Education," *Quarterly Journal of Economics* 114:2 (1999), 577–599. 10.1162/003355399556070
- Bursztyn, L., B. Ferman, S. Fiorin, M. Kanz, and G. Rao, "Status Goods: Experimental Evidence from Platinum Credit Cards," NBER working paper 23414 (2017).
- Bursztyn, L., T. Fujiwara, and A. Pallais, "Acting Wife: Marriage Market Incentives and Labor Market Investments," *American Economic Review* 107:11 (2017), 3288–3319. 10.1257/aer.20170029
- Cavallo, A., G. Cruces, and R. Perez-Truglia, "Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Experiments," *American Economic Journal: Macroeconomics* 9:3 (2017), 1–35. 10.1257/mac.20150147
- Clark, A. E., and A. J. Oswald, "Satisfaction and Comparison Income," *Journal of Public Economics* 61:3 (1996), 359–381. 10.1016/0047-2727(95)01564-7
- Clark, A., C. Senik, and K. Yamada, "When Experienced and Decision Utility Concur: The Case of Income Comparisons," *Journal of Behavioral and Experimental Economics* 70 (2017) 1–9. 10.1016/j.socec.2017.07.002
- Clark, A. E., N. Westergård-Nielsen, and N. Kristensen, "Economic Satisfaction and Income Rank in Small Neighbourhoods," *Journal of the European Economic Association* 7:2–3 (2009), 519–527. 10.1162/JEEA.2009.7.2-3.519
- Couture, V., and J. Handbury, "Urban Revival in America 2000 to 2010," University of Pennsylvania mimeograph (2016).
- Cruces, G., R. Perez-Truglia, and M. Tetaz, "Biased Perceptions of Income Distribution and Preferences for Redistribution: Evidence from a Survey Experiment," *Journal of Public Economics* 98 (2013), 100–112. 10.1016/j.jpubeco.2012.10.009
- Daniel, C., and M. O'Brien, "Why Study Medicine?" *Student Doctor Network* (April 24, 2008).
- Doob, A. N., and A. E. Gross, "Status of Frustrator as an Inhibitor of Horn-Honking Responses," *Journal of Social Psychology* 76:2 (1968), 213–218. 10.1080/00224545.1968.9933615
- Easterlin, R. A., "Does Economic Growth Improve the Human Lot? Some Empirical Evidence," *Nations and Households in Economic Growth* 89 (1974), 89–125.
- Fennis, B. M., "Branded into Submission: Brand Attributes and Hierarchization Behavior in Same-Sex and Mixed-Sex Dyads," *Journal of Applied Social Psychology* 38:8 (2008), 1993–2009. 10.1111/j.1559-1816.2008.00377.x
- Ferrer-i-Carbonell, A., "Income and Well-Being: An Empirical Analysis of the Comparison Income Effect," *Journal of Public Economics* 89:5 (2005), 997–1019. 10.1016/j.jpubeco.2004.06.003
- Festinger, L., "A Theory of Social Comparison Processes," *Human Relations* 7:2 (1954), 117–140. 10.1177/001872675400700202
- Fisman, R., S. S. Iyengar, E. Kamenica, and I. Simonson, "Gender Differences in Mate Selection: Evidence from a Speed Dating Experiment," *Quarterly Journal of Economics* 121:2 (2006), 673–697. 10.1162/qjec.2006.121.2.673
- Frank, R. H., *Choosing the Right Pond: Human Behavior and the Quest for Status* (New York: Oxford University Press, 1985a).
- , "The Demand for Unobservable and Other Nonpositional Goods," *American Economic Review* 75:1 (1985b), 101–116.
- Gautier, P. A., M. Svarer, and C. N. Teulings, "Marriage and the City: Search Frictions and Sorting of Singles," *Journal of Urban Economics* 67:2 (2010), 206–218. 10.1016/j.jue.2009.08.007
- Hauser, O. P., and M. I. Norton, "(Mis)perceptions of Inequality," *Current Opinion in Psychology* 18 (2017), 21–25. 10.1016/j.copsyc.2017.07.024, PubMed: 29221507
- Heffetz, O., "A Test of Conspicuous Consumption: Visibility and Income Elasticities," this REVIEW 93:4 (2011), 1101–1117.
- Hitsch, G., A. Hortaçsu, and D. Ariely, "What Makes You Click: An Empirical Analysis of Online Dating," *Quantitative Marketing and Economics* 4 (2010), 393–427. 10.1007/s11129-010-9088-6
- Johansson-Stenman, O., F. Carlsson, and D. Daruvala, "Measuring Future Grandparents' Preferences for Equality and Relative Standing," *Economic Journal* 112:479 (2002), 362–383. 10.1111/1468-0297.00040
- Kahneman, D., and A. Tversky, "Subjective Probability: A Judgment of Representativeness," *Cognitive Psychology* 3:3 (1972), 430–454. 10.1016/0010-0285(72)90016-3
- Karadja, M., J. Mollerstrom, and D. Seim, "Richer (and Holier) Than Thou? The Effect of Relative Income Improvements on Demand for Redistribution," this REVIEW 99:2 (2017), 201–212.
- Kuziemko, I., R. W. Buell, T. Reich, and M. I. Norton, "Last-Place Aversion: Evidence and Redistributive Implications," *Quarterly Journal of Economics* 129:1 (2014), 105–149. 10.1093/qje/qjt035
- Li, S., "Obviously Strategy-Proof Mechanisms," *American Economic Review* 107:11 (2017), 3257–3287. 10.1257/aer.20160425
- Luttmer, E. F., "Neighbors as Negatives: Relative Earnings and Well-Being," *Quarterly Journal of Economics* 120:3 (2005), 963–1002.

- Nelissen, R. M., and M. H. Meijers, "Social Benefits of Luxury Brands as Costly Signals of Wealth and Status," *Evolution and Human Behavior* 32:5 (2011), 343–355. 10.1016/j.evolhumbehav.2010.12.002
- National Resident Matching Program, "Results of the 2015 NRMP Applicant Survey by Preferred Specialty and Applicant Type" (Washington, DC: National Resident Matching Program, 2015).
- Perez-Truglia, R., "The Effects of Income Transparency on Well-Being: Evidence from a Natural Experiment," mimeograph (2016).
- Political Conformity: Event-Study Evidence from the United States," *this REVIEW* 100:1 (2018), 14–28.
- Perez-Truglia, R., and G. Cruces, "Partisan Interactions: Evidence from a Field Experiment in the United States," *Journal of Political Economy* 125:4 (2017), 1208–1243. 10.1086/692711
- Ray, D., and A. Robson, "Status, Intertemporal Choice, and Risk-Taking," *Econometrica* 80:4 (2012), 1505–1531. 10.3982/ECTA9042
- Rees-Jones, A., "Suboptimal Behavior in Strategy-Proof Mechanisms: Evidence from the Residency Match," *Games and Economic Behavior* 108C (2018), 317–330. 10.1016/j.geb.2017.04.011
- Rees-Jones, A., and S. Skowronek, "Why Do We Lie in Incentive-Compatible Mechanisms? Evidence from the Residency Match," mimeograph (2017).
- Roback, J., "Wages, Rents, and the Quality of Life," *Journal of Political Economy* 90:6 (1982), 1257–1278. 10.1086/261120
- Rosen, S., "Wage-Based Indexes of Urban Quality of Life" (pp. 74–104), in P. Mieszkowski and M. Straszheim, eds., *Current Issues in Urban Economics* (Baltimore: Johns Hopkins University Press, 1979).
- Roth, A. E., and E. Peranson, "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design," *American Economic Review* 89:4 (1999), 748. 10.1257/aer.89.4.748, PubMed: 29115787
- Senik, C., "When Information Dominates Comparison: Learning from Russian Subjective Panel Data," *Journal of Public Economics* 88:9 (2004), 2099–2123. 10.1016/S0047-2727(03)00066-5
- Shigeoka, H., and K. Yamada, "Income-Comparison Attitudes in the US and the UK: Evidence from Discrete-Choice Experiments," NBER working paper 21998 (2016).
- Sledge, W. H., P. Leaf, and M. Sacks, "Applicants' Perceptions of Psychiatric Residency Training Programs," *Academic Psychiatry* 13:1 (1989), 24–30. 10.1007/BF03341837, PubMed: 24442973
- Solnick, S. J., and D. Hemenway, "Is More Always Better? A Survey on Positional Concerns," *Journal of Economic Behavior and Organization* 37:3 (1998), 373–383. 10.1016/S0167-2681(98)00089-4
- Tideman, S., P. Frijters, and M. A. Shields, "Relative Income, Happiness, and Utility: An Explanation for the Easterlin Paradox and Other Puzzles," *Journal of Economic Literature* 46:1 (2008), 95–144. 10.1257/jel.46.1.95
- Van Gelder, M. M., R. W. Bretveld, and N. Roeleveld, "Web-Based Questionnaires: The Future in Epidemiology?" *American Journal of Epidemiology* 172:11 (2010), 1292–1298. 10.1093/aje/kwq291, PubMed: 20880962
- Yamada, K., and M. Sato, *Another Avenue for Anatomy of Income Comparisons: Evidence from Hypothetical Choice Experiments* (Berlin: Springer, 2016).